U.S. Department of Energy
Office of the Chief Financial Officer

Technical Architecture Plan
for
iManage Data Warehouse/iPortal

February 2012
Version 2.3 (modified)

Draft

Title Page

Document Name: Technical Architecture Plan for the iManage Data Warehouse/Portal (IDW/iPortal)

Publication Date: February 2012

Draft

Technical Architecture Plan

1. Introduction

1.1 Background

1.1.1 The Integrated Management Navigation System (iManage) Program is the Department's solution for managing enterprise-wide systems initiatives to achieve improved financial and business efficiencies, integrated budget and performance, and expanded electronic government in support of the President's Management Agenda. The iManage Program is a collaborative Departmental effort to define and provide a modern, integrated corporate business system for the Department of Energy. The Project Portfolio is comprised of enterprise-wide systems initiatives to include: the Standard Accounting and Reporting System (STARS), IDW/iPortal, iManage iBudget, Strategic Integrated Procurement Enterprise System (STRIPES), Corporate Human Resource Information System (CHRIS), and the E-Travel System (eTS).

As a key component of the iManage Program, the Department will rely on IDW/iPortal for internal executive, management, and operational reporting, and as the authoritative source for external ad-hoc requests for data. This data-centric approach to managing and integrating data will allow the Department to rapidly respond to new and changing demands for information. This effort is important not only to future system development, but also to legacy systems by creating the ability to integrate information from all organizational elements in a single corporate repository.

1.2 Purpose of this Document

1.2.1 The purpose of this IDW/iPortal Technical Architecture Plan is to:

1.2.1.1 Increase understanding of data warehousing terms and concepts.

1.2.1.2 Provide an overview of the data warehouse services, products, and infrastructure that will collectively define the data warehouse architecture.

1.2.1.3 Document the approach used by the IDW/iPortal Team in selecting an appropriate architecture for the data warehouse.

1.2.1.4 Provide an overview of iPortal services provided.

1.2.1.5 Document the approach used by IDW/iPortal team in selecting the appropriate design for the portal.

1.3 Objectives

1.3.1 This document is organized into five sections that address the goals stated in the purpose above.

1.3.1.1 Section 1 – Introduction. This section provides an introduction to the document. It describes the objectives of the IDW/iPortal initiative.

1.3.1.2 Section 2 – Major Architectural Elements. Major Architectural Elements. This section provides an overview of the fundamental data warehouse data store, service, and feature concepts that provide a framework for addressing the IDW/iPortal architecture.

1.3.1.3 Section 3 – High Level ETL Process Framework. This section provides overview of key elements of IDW/iPortal ETL process.

1.3.1.4 Section 4 – Oracle Data Warehouse Products. This section provides an overview of Oracle data warehouse products.

2. Major Architectural Elements

2.1 High-level Technical Architecture Model.

2.1.1 The figure below illustrates the technical architecture model for the data warehouse. This is commonly divided into two pieces: The Back Room and The Front Room.

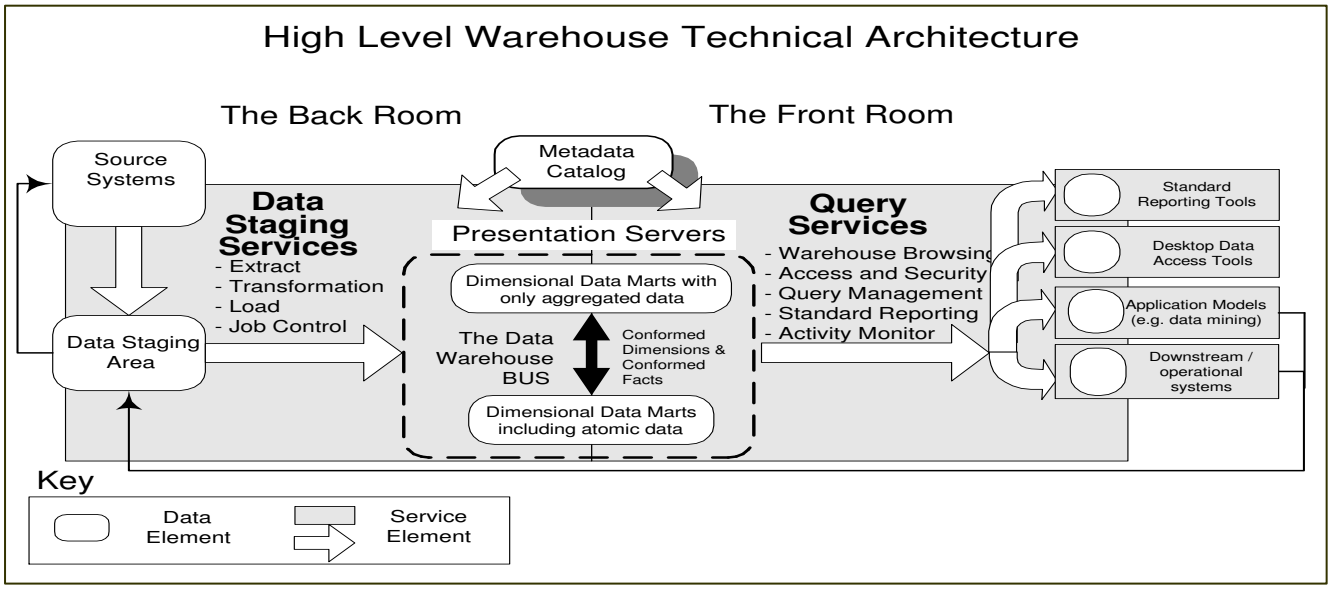

2.2 Back Room.

2.2.1 The Back Room is the "engine room" of the data warehouse. Here Data Architect, Database Administrator, and System Administrators move data from Source Systems to Presentation Servers using job scheduled and monitored execution of extraction, transformation and loading processes, known as ETL, or "Data Staging Services".

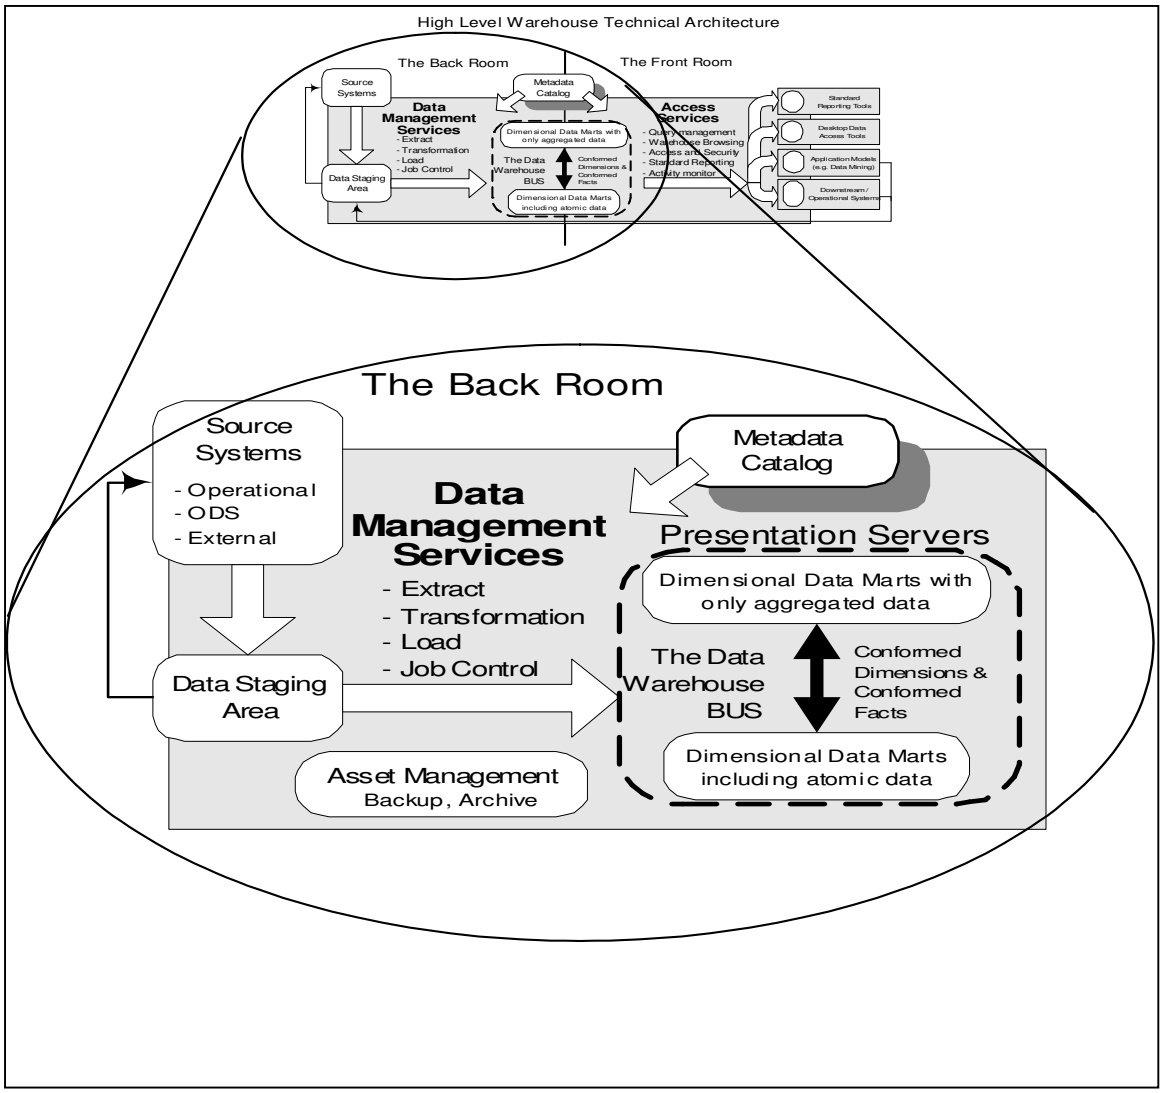
High Level Warehouse Technical Architecture
The Back Room
The Front Room
Source Systems
Metadata Catalog
Data Management Services
- Extract
- Transformation
- Load
- Job Control
Data Staging Area
Dimensional Data Marts with only aggregated data
The Data Warehouse BUS
Conformed Dimensions & Conformed Facts
Dimensional Data Marts including atomic data
Access Services
- Query management
- Warehouse Browsing
- Access and Security
- Standard Reporting
- Activity monitor
Standard Reporting Tools
Desktop Data Access Tools
Application Models (e.g. Data Mining)
Downstream / Operational Systems
The Back Room
Source Systems
- Operational
- ODS
- External
Metadata Catalog
Data Management Services
- Extract
- Transformation
- Load
- Job Control
Presentation Servers
Dimensional Data Marts with only aggregated data
The Data Warehouse BUS
Conformed Dimensions & Conformed Facts
Dimensional Data Marts including atomic data
Data Staging Area
Asset Management
Backup, Archive

2.2.2 Back Room Services

2.2.2.1 Data Staging Area Services. Data staging area services are tools and techniques employed in the data staging process. The services provided by the staging area are similar to services provided by software development tools. A listing of the services is contained in the table below:

Data Staging Services
Code library management check in/check out
Version control
Production and development system builds
Usability
GUI
System documentation
Metadata driven

2.2.2.2 Metadata Catalog Maintenance Services. Metadata plays a large role in the data warehouse environment. Metadata serves as documentation for the contents and processes of the data warehouse. To that end services are needed to create, capture, manage, and report on metadata. A listing of the services is contained in the table below:

Metadata Catalog Maintenance Services
Information catalog integration/merge (e.g., from the data model to the database to the front-end tool)
Metadata management (e.g., remove old, unused entries)
Capture existing metadata (e.g., DDL from mainframe or other sources)
Manage and display graphical or tabular representations of the metadata catalog contents (the metadata browser)
Maintain user profiles for application and security use
Security for the metadata catalog
Local or centralized metadata catalog support

Metadata information on tables, columns, and jobs are needed to create and maintain the warehouse rather than hard coding this information directly into SQL code. A listing of useful metadata driven data staging procedures are contained in the table below:

Useful Metadata-Driven Data Staging Procedures
Create table
Create index
Drop table
Drop index
Copy table
Add user
Drop user

Useful Metadata-Driven Data Staging Procedures
Run SQL Script
Run DBMS Table Loader

2.2.2.3 Extract Services. Extracting data from the source systems is probably the largest single effort in the data warehouse project. Extraction services allow for any data source to readily be captured, and extracted into the data staging area for transformation and then into the data warehouse target presentation area. A listing of the services are contained in the table below:

Extract Services
Support for multiple sources
Code generation
Multiple extract types
Incremental loads
Transaction events
Full refresh
Replication
Compression/decompression

2.2.2.4 Data Transformation Services. After the data extraction phase, various tasks are performed to convert source data into a useable form for analytic business reporting. Steps like re-mapping activity codes into new codes, conforming lookup table values, removing duplicates are examples of transformations. Below are some of the transformation services likely to be required in a data warehouse.

Data Transformation Services
Integration
Slowly changing dimension maintenance
Referential integrity checking
Denormalization and renormalization
Cleansing, deduping, merge/purge
Data type conversion
Calculation, derivation, allocation
Aggregation
Data content audit
Data lineage audit
Tool- or analysis-specific transformation
Null value handling
Pre- and post-step exits

2.2.2.5 Data Loading Services. The capabilities needed during the data loading process are primarily a function of the target platform. Some of these capabilities are listed in the table below:

Data Loading Services

Data Loading Services
Support for multiple targets
Load optimization
Entire load process support

2.2.2.6 Data Staging Job Control Services. The entire data staging job stream should be managed, to the extent possible, through a single, metadata-driven job control environment. The job control process should also capture metadata regarding the progress and statistics of the daily job itself. A listing of job control services are provided below:

Data Staging Job Control Services
Job definition
Job scheduling
Monitoring
Logging
Exception handling
Error handling
Notification

2.2.2.7 Back Room Asset Management Services. The data warehouse is subject to risks of data loss just as any other computer system. In addition to these risks, the data warehouse also needs to keep more data for longer periods of time than operational systems. The services needed to assist the data warehouse get back online after system interruptions are listed below:

Back Room Asset Management Services
Backup and Recovery
Archive and Retrieval

2.2.3 Back Room Data Stores

2.2.3.1 Landing Area. Data is extracted from the source operational systems and transported to the landing area. The landing area servers as a local copy of source data for use in the downstream ETL processes.

2.2.3.2 Staging Area. Once data has been transported to the landing area it is available for the transform process. During the transform process various tasks are performed to convert source data resident the landing area into a useable form for analytic business reporting. Steps like re-mapping activity codes into new codes, conforming lookup table values, and removing duplicates are examples of transformations. After data has been integrated with other data, cleansed, and transformed the data is then stored in a staging area. The staging area is a holding area for data that will eventually be transported via the load process into the target data warehouse.

2.3 The Front Room

2.3.1 The front room is the public face of the data warehouse, providing a layer between the users and the information that hides the complexities of the data. Front Room services consist of browsing, security, monitoring, reporting, administration, and querying.

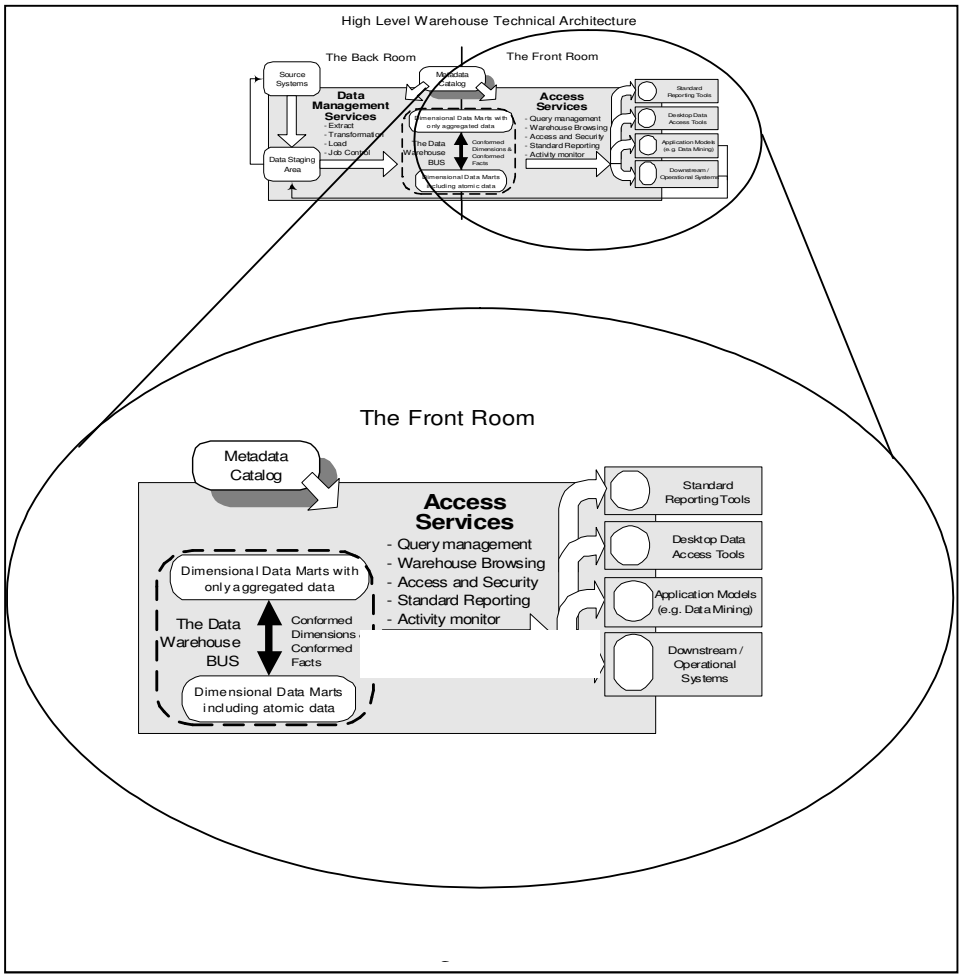

2.3.2 Front Room Services

2.3.2.1 Warehouse Browsing Services. Warehouse browsing takes advantage of the metadata catalog to support the users in their efforts to find and access the information they need. The warehouse browser should be dynamically linked to the metadata catalog to display currently available subject areas and the data elements within those subject areas.

Warehouse Browsing Services
Subject areas
Data elements
Definitions
Derivations
Standard reports with elements
Link to canned report, tool or report scheduler

2.3.2.2 Access and Security Services. Access and security services facilitate a user's connection to the data warehouse.

2.3.2.3 Activity Monitoring Services. Activity monitoring services capture information about the use of the data warehouse.

Activity Monitoring Services
Performance
User support
Marketing
Planning

2.3.2.4 Query Management Services. Query management services are the set of capabilities that manage the exchange between the query information, the execution of the query on the database, and the return of the result set to the desktop.

Query Services
Metadata driven
Content simplification
Query reformulation
Query retargeting and multipass SQL
Aggregate awareness
Date awareness
Query governing

2.3.2.5 Standard Reporting Services. Standard reporting provides the ability to create production style fixed format reports that allow user interaction, support a broad audience, and regular execution schedules.

Standard Reporting Services
Report development environment

Standard Reporting Services
Report execution server
Parameter or variable driven capabilities
Time- and event-based scheduling of report execution
Iterative execution
Flexible report definitions
Flexible report delivery
Via multiple delivery methods (e-mail, Web, network directory, desktop directory)
Multiple result types (data access tool file, database table, spreadsheet)
User accessible publish and subscribe
Report linking
Report library with browsing capability
Mass distribution
Report environment administration tools

2.3.2.6 Desktop Services. Desktop services are found in front-end tools that provide users access to the data warehouse. Desktop services need to support multiple consumer types each with varying technical skill levels and degrees of analytical sophistication.

- The table below details characteristics of different consumers of data warehouse information:

Usage Area	User Type			
	Paper User	Executive Information System Access	Simple Ad-Hoc	Power User
Data Warehouse	Rely on others to Navigate	Standard reports, default parameters, EIS (Executive Information system)	Create simple queries, modify existing queries, browse/change parameters, navigate hierarchies	Build full queries from scratch, direct database access

- The table below maps information needs and data access categories to user roles and Oracle products:

Info	Needs Category	User Roles	Data Access Category	Common Tools	Audience Size
High level monitoring	Key metrics, flags	Senior Management	Executive Information System Access "dashboard"	Oracle Financial Intelligence, Oracle Discoverer Viewer, BI Suites	Small
Business Tracking	Budgets, funds, accounts, etc; drill down to detail	Mid-Management and business managers	Standard reports, parameter driven	Oracle Discoverer Viewer, BI Publisher	Large

Investigating	Exceptions; new problems or opportunities; business case development	Same as above plus business analysts	Ad-hoc analysis	Discoverer Plus, BI Suites	Medium
Complex Analysis	Composite querying, statistical analysis, model development	Business analysts and Analytical experts	Data mining, Advanced Analysis	Discoverer Plus, Oracle Express / Oracle OLAP, BI Suites	Small

2.3.2.7 Data Access Tools. The quality of the user's overall experience is heavy influenced by the capabilities provided by the tools chosen to provide access to the data warehouse. A table of common data access tool capabilities is provided below.

Data Access Tool Capabilities
Executive Information System Access
Easy generation of sophisticated, interactive, engaging front-end screens
Simple navigation controls
Automatic replacement or on-the-fly creation of underlying report contents
User interface controls for conceptual representations like stoplights (high, medium, low), gauges, and sophisticated charts
Geographical charts with links to underlying reports
"Alerters" or exception controls that monitor specified values, ranges, or differences and notify the user when they exceed target levels
Ability to define and interact with multiple simultaneous connections to multiple data sources
Standard reports
Ad Hoc Capabilities
Query formulation
Multipass SQL
Highlighting
Successive constraints
Semiadditive summations
ANSI SQL 92 support
Direct SQL entry
Analysis and presentation capabilities
Basic calculations on the results set
Pivot the results
Column calculations on pivot results
Column and row calculations
Sorting
Complex formatting
Charting and graphs
User-changeable variables

Data Access Tool Capabilities
User Interaction
Ease of use
Metadata access
Pick lists
Seamless integration with other applications
Export to multiple file types, including HTML
Embedded queries
Technical features
Multitasking
Cancel query
Scripting
Connectivity
Scheduling
Metadata driven
Software administration
Security
Querying

3. High Level ETL Process Framework

3.1 ETL Process Flow

Figure below illustrates the redesigned ETL process flow from a high level perspective. Key to the new design for the ETL is that it be as close to real-time as possible as well as non-intrusive. The resource intensive activities will continue to be conducted in stage area making them transparent to the users in the target or "live" area of the warehouse. The redesign takes advantage of the latest features in Oracle to meet the design objective set forth by IDW/iPortal management.

HIGH LEVEL ETL FLOW

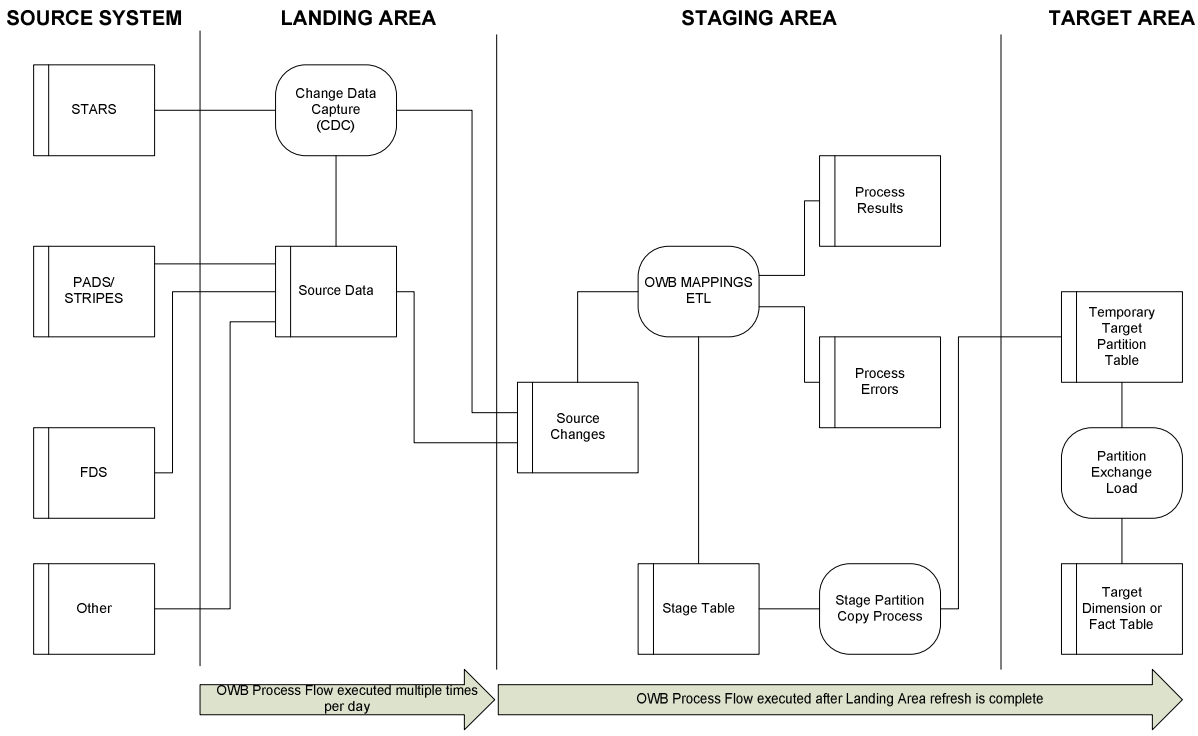

3.1.1 Landing Area

The mechanisms used to load the data into the landing area from the STARS source will be altered as part of the redesign to utilize the change data capture (CDC) features in the Oracle database. The load processes for non-STARS sources will remain unchanged. Also, the landing area load processes will no longer be scheduled and executed from the UNIX crontab. The load processes will be scheduled and executed via a process flow from within OWB.

As data is changed in STARS, CDC will write the changed records to a change log that resides in the warehouse. The new ETL process is designed to read the change logs maintained by CDC performing an incremental load rather than a full refresh. Without CDC identifying changed records, incremental loading would not be possible.

A detailed discussion of the planned IDW/iPortal CDC process is included in a document titled "IDW Incremental Load Strategy and Materialized View Usage Criteria".

3.1.2 Staging Area

Processing within the staging area will be completely redesigned for incremental processing and can be executed during normal business hours in the background. None of the existing ETL processes can be used with this design since the paradigm used for loading is completely different. Incremental loading satisfies the "real-time" design objective by only processing what has changed on the source system.

The ETL processes will read from tables containing the changed records identified by the CDC process and load them into the stage tables. In addition, the ETL transformation logic will be incorporated into the mappings rather than views as was previously done facilitating ease of maintenance. A set of new standard pre and post mapping routines will be developed for inclusion in each of the ETL maps. These routines will maintain IDW/iPortal process tables used to record detailed ETL status information for use by the operations team.

The IDW/iPortal process tables will provide the operations team with visibility as to the status of the ETL process at any point in time. In addition, the IDW/iPortal process data may be used for trending and performance analysis. The tables containing the ETL results will grow large over time and the operations team will need to decide when to archive the data.

3.1.3 Target Area

When the stage table loads are complete, a new process will identify the partitions that have changed in stage and exchange those with the ones in the target area. The partition exchange feature in Oracle will be used to facilitate the new target load method. This feature is explained in greater detail later in this document.

The exchange is extremely fast and can occur while the target area is in use by the IDW/iPortal users satisfying the non-intrusive design objective. This technique is presently employed within the IDW/iPortal for loading the custom balances table from STARS into the landing area so it is proven and the IDW/iPortal DBA resources are familiar with its use.

3.2 Change Data Capture (CDC)

Figure below illustrates the CDC process flow from a high level perspective. The proposed CDC process is described in detail in the document entitled "Incremental Load Strategy and Materialized View Usage Criteria".

HIGH LEVEL CDC FLOW

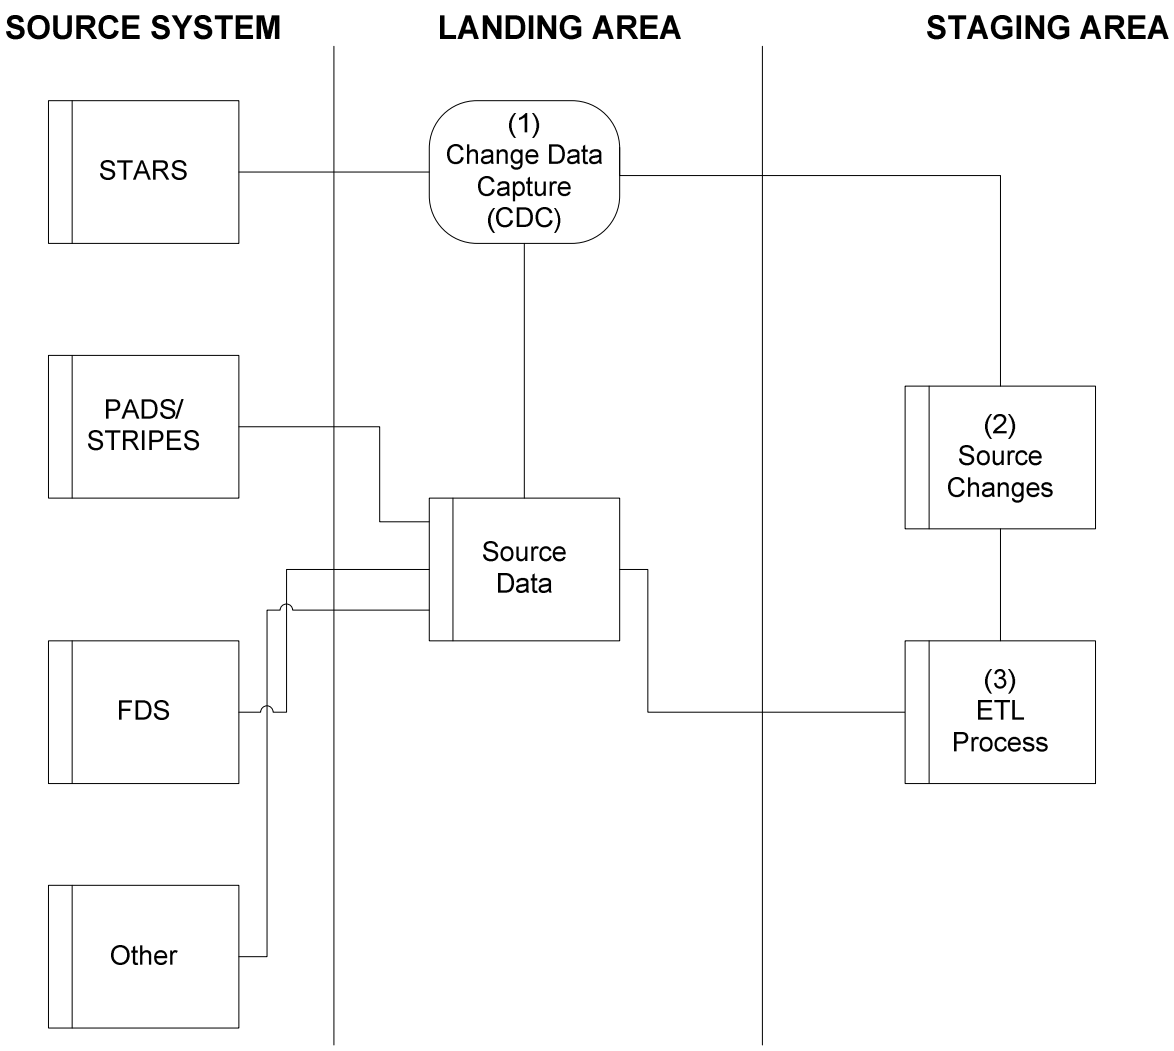

The steps in the proposed CDC process are as follows and are illustrated in Figure 3 above.

1) CDC captures the data manipulation occurring on each source table as it occurs on the source
2) The inserts, deletes and "true" record changes since the last ETL execution are published to the stage area.
3) The redesigned ETL will process the changes

3.3 Partition Exchange Loading (PEL)

The partition exchange feature of Oracle swaps definitions of named partitions and a named table so that the data suddenly exists in the right place in the partitioned table. In addition, there are two optional clauses allowing indexes to be swapped as well. Oracle will not check whether the data actually belongs in the partition so the exchange is very quick.

One possible draw back is that two users executing the same query at slightly different times will see different results. This can occur if the under lying partition that the query is reading is in the process of being exchanged.

To utilize this feature all dimension and fact table will be range partitioned in IDW/iPortal. Smaller dimension tables will likely be housed in a single partition.

Key features of PEL are as follows:

- "In-flight" queries continue to process during and after EXCHANGE PARTITION command operations
- Oracle's read-consistency mechanisms cause existing operations to use data that was exchanged away from the table, and new operations to use data exchanged into the table
- Local-partitioned indexes and statistics are exchanged as well
- Global-partitioned indexes are maintained during exchange operation
- The basic technique of bulk-loading new data into a temporary "load table", which is indexed, analyzed, and then "published" all at once to end-users using the EXCHANGE PARTITION operation, will be the default load technique for all IDW/iPortal target tables
- Exchange Partition load techniques make "load windows" of restricted activity unnecessary

Figure below depicts the planned PEL flow at a high level.

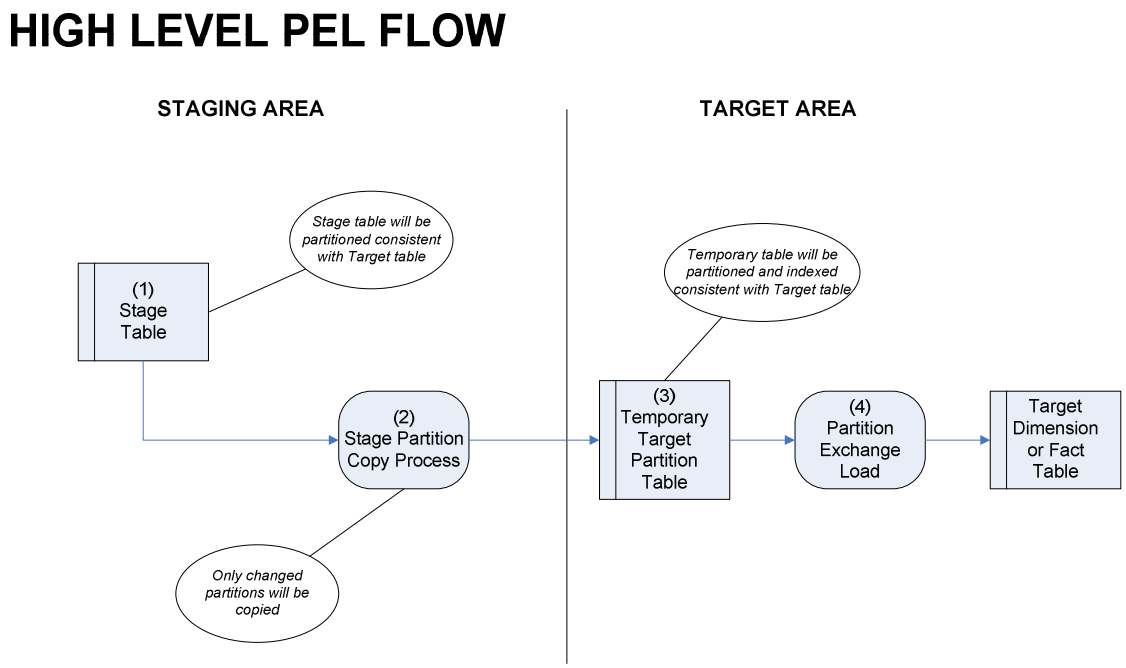

The steps in the proposed PEL process are as follows and are illustrated in Figure 4 above.

1. The stage table is loaded by the ETL.
2. A new copy process is employed to identify changed table partitions and copy them to a temporary table in the target area. Note, this temporary table is identical in structure to the target table partition it is intending to replace.
3. The temporary table is loaded, indexed and analyzed.

The partition exchange command is issued to load the new data into the partition of the target table.

3.4 OWB Process Flows

The "Oracle Warehouse Builder User's Guide" describes a process flow as follows:

"A process flow describes dependencies between Warehouse Builder mappings and external activities such as email, FTP, and operating system commands."

Presently, the IDW/iPortal does not make use of this facility in OWB however it will be used extensively in the redesigned IDW/iPortal. Every effort will be made

to replace the existing UNIX scripts with OWB process flows.

Each of the IDW/iPortal process flows will have the same flow at a high level as illustrated in figure below. The key design element is the testing of incompatibilities to ensure that no incompatible process is running when the process attempts to run. By its nature, process flows are able to link together tasks and; therefore, initiate processes in a defined sequence. However, it will not prevent you from running a given process repeatedly or with another incompatible process. The new IDW/iPortal process tables are not intended to replace the functionality found in the OWB process flow but to augment it.

HIGH LEVEL OWB PROCESS FLOW

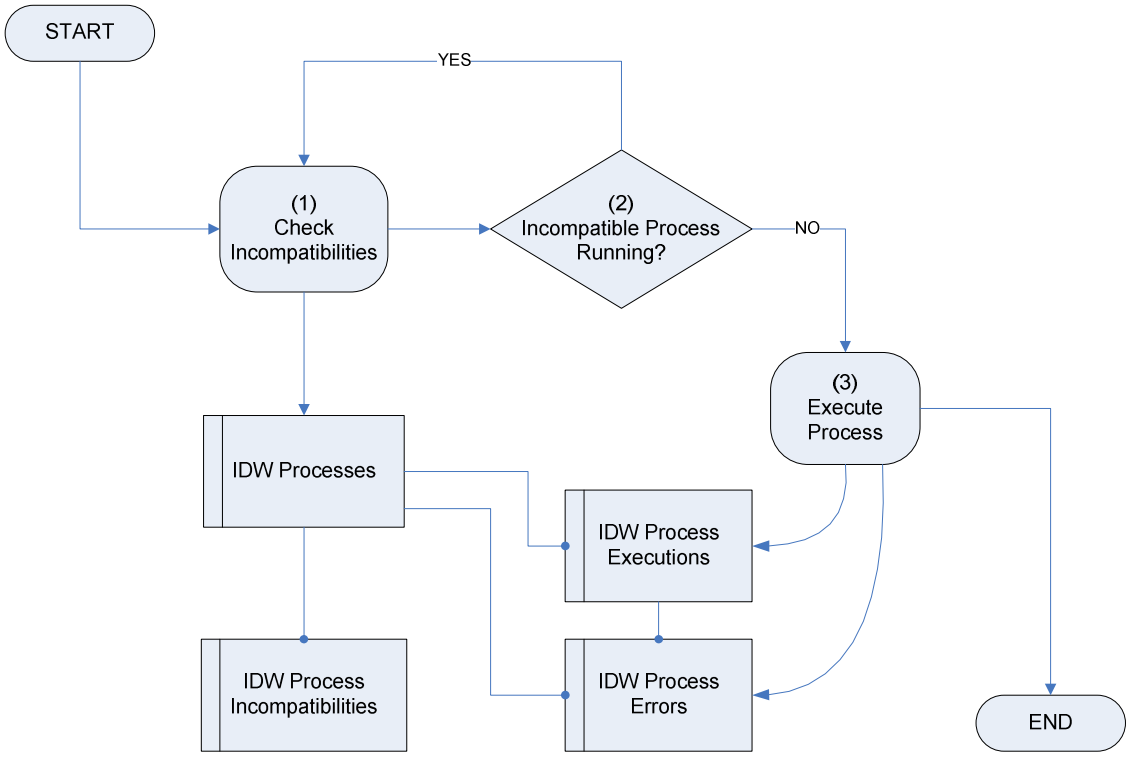

The steps in the common to all OWB process flows in IDW/iPortal are as follows and are illustrated in Figure 5 above.

1. The IDW/iPortal Process Incompatibilities table will be read at the beginning of each process
2. If an incompatible process is running then the new process will wait for (n) minutes loop back and check for incompatibilities again
3. If no incompatible processes are running then the process will execute and record its results in the IDW/iPortal process tables

4. IDW Software Products

4.1 Overview

4.1.1 Beginning in late 2005, Oracle released the Oracle 10g line of its core technology products: the Oracle 10g R2 Database; Oracle Fusion Middleware Family; Oracle 10g R2 Application Server; 10g R2 Warehouse Builder and Oracle Developer Suite. Components previously available separately have now been integrated into suites.

4.2 Oracle Enterprise Database.

Many features were added to the Oracle 11g Database to improve the performance, manageability, and scalability of the data warehouse. These features included the following:

4.2.1 Enhanced index performance and manageability.

4.2.2 An increased number of data and index partitioning schemes.

4.2.3 Improvements supporting the ETL processes:

4.2.3.1 Change data capture provides a mechanism to identify changed rows in an Oracle database

4.2.3.2 External tables that allow transformations to be executed while loading the data in the database.

4.2.3.3 Table functions that provide pipelined parallel execution of PL/SQL transformations.

4.2.4 Materialized views that can be incrementally refreshed after individual DML statements, as well as after direct path loads.

4.2.5 On-line performance tuning of Oracle SGA.

4.2.6 On-line object maintenance, i.e. recreate index, move partition, redefine object, etc.

4.3 Oracle Enterprise Manager 11g.

Spanning applications, middleware, and database management, Oracle Enterprise Manager 11g Release 3's unique "top-down" approach enables Information Technology (IT) departments to focus on what matters to the business - greater agility, better service quality and lower operational costs. Oracle Enterprise Manager enables its customers to manage their applications from top to bottom - from monitoring service levels to proactively isolating business exceptions before they escalate, and re-mediating issues at any level of the IT stack.

OEM is designed to run in a three-tier architecture, as follows:

- The console provides the graphical user interface on the client systems.
- The management servers (OMS) provide administrative functions, such as executing jobs and events, and runs in the middle tier.

- The intelligent agents monitor the database, start and stop it, and gather performance data. The intelligent agents run in the third tier and are located on the same nodes as the Oracle databases.

4.4 Oracle Designer

4.4.1 Oracle Designer is a software tool for analyzing business requirements and for designing and generating client/server systems that meet those requirements. Oracle Designer incorporates support for business process modeling, systems analysis, software design, and system generation.

4.4.2 Oracle Designer provides a multi-user repository and is closely integrated with Oracle Warehouse Builder and Developer, Oracle's client/server development toolset. In this way, Oracle Designer allows organizations to design and rapidly deliver scalable, client/server systems that can adapt to changing business needs.

4.4.3 The Oracle Designer tools are grouped into areas that reflect the primary needs of its different types of user:

- Modeling System Requirements. These tools are used to: model business processes; re-examine the methods used to achieve the goals of the organization; create diagrammatic representations of business processes; record details about them; describe business requirements in detail; create logical data models, functions and flows of data in the systems that make up the organization.
- Transforming Preliminary Designs. The transformer tool is used to generate physical implementations of logical data models.
- Designing and Generating. The Design Editor is used to: design a system that meets the business requirements of an organization; provide a development environment for system engineers and designers; create server-side components and client-side applications from definitions recorded in the Repository.
- Repository Tools. These tools are used to: enter and edit information in the Repository; show relationships between elements in the Repository; generate predefined and custom reports on the contents of the Repository; administer Repositories.

Oracle Warehouse Builder

Oracle10g Warehouse Builder is a data warehousing tool that enables one to design a complete logical model of the warehouse. Warehouse Builder also integrates with Oracle database and query tools for a complete business intelligence product.

- Integration with Oracle Portal provides reports on metadata, including lineage and impact analysis.
- Warehouse Builder integrates with Business Intelligence tools, such as Oracle Discoverer, for reporting, analysis, and data mining capability. It can also integrate application data and metadata from other sources.

Oracle10g Warehouse Builder is a data warehousing tool used to:

- Design. The Warehouse Builder user interface contains graphical editors that enable one to design a complete logical model of the warehouse.
- Extract, transform, and load. Warehouse Builder includes a graphical interface that enables one to plan how to extract data from a variety of sources, transform it, and configure it for loading into the data warehouse.
- Deploy and Update. Warehouse Builder code generator enables one to deploy and populate a data warehouse or data mart without manual coding.
- Manage. Warehouse Builder integrates with Oracle Workflow and Oracle Enterprise manager, enabling one to schedule data warehouse updates.
- Integrate. Warehouse Builder integrates with Oracle database and query tools for a complete business intelligence product.

Warehouse Builder consists of the following parts:

- Repository: A set of relational objects in an Oracle 10g database that store metadata definitions used in the data warehouse. The repository also stores transformation libraries.
- Client: A graphical user interface that enables one to interact with the repository using wizards and editors. The client is used to logically design the data warehouse and generate the code scripts used for deploying and loading a data warehouse.
- Runtime Environment: A set of Oracle objects that one use to set up the target schema. It provides auditing capabilities and a graphical user interface for viewing audit messages.
- Browser: Integration with Oracle Portal provides reports on the metadata, including lineage and impact analysis.
- Application Integrators: Warehouse Builder integrates with Business Intelligence tools, such as Oracle Discoverer, for reporting, analysis, and data mining capability. One can also integrate application data and metadata from sources such as SAP.

Oracle Fusion Middleware.

- Oracle Fusion Middleware is a portfolio of leading, standards-based and customer-proven software products that spans a range of tools and services from J2EE and developer tools, to integration services, business intelligence, collaboration, and content management. This comprehensive and well-integrated family of products offers complete support for development, deployment, and management of Service-Oriented Architecture. Oracle Fusion Middleware's hot-pluggable architecture allows you to leverage your investments in any existing application, system, or technology. And because maximum uptime is good for business, the unbreakable qualities of Oracle Fusion Middleware's core technology minimize the disruption of planned or unplanned outages.

Oracle Application Server

- Oracle Application Server 10g offers a comprehensive solution for developing, integrating, and deploying your enterprise's applications, portals, and Web services. Based on a powerful and scalable J2EE server, Oracle Application Server 10g provides complete business integration and business intelligence suites, and best-of-breed portal software. As the only platform designed for grid computing as well as full lifecycle support for Service-Oriented Architecture (SOA), Oracle Application Server gives you unmatched scalability, availability, manageability, and security.

Oracle Business Intelligence Suite Enterprise Edition Plus

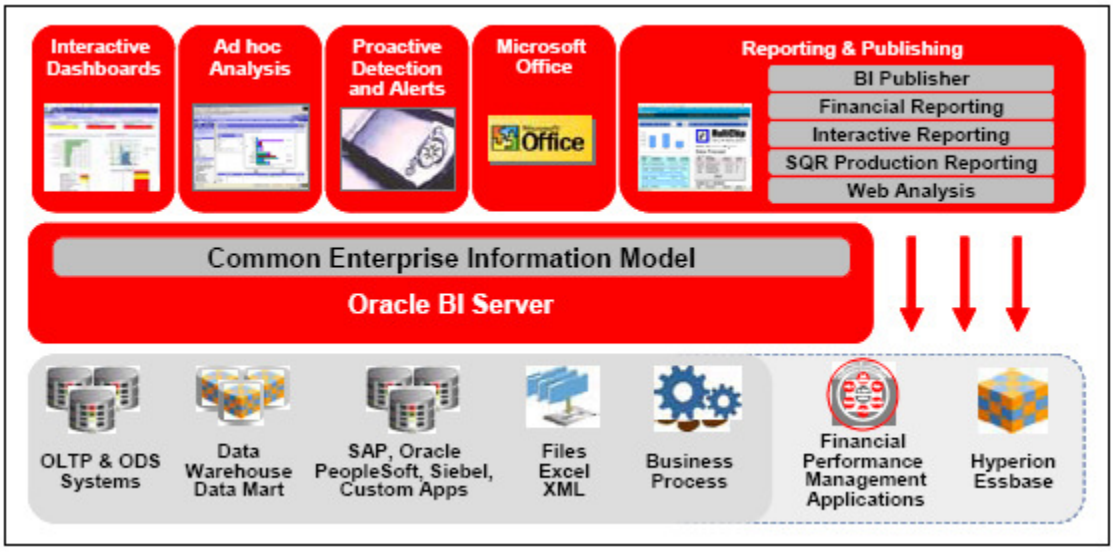

- Oracle BI EE Plus is a comprehensive suite of enterprise BI products that delivers a full range of analysis and reporting capabilities. Featuring a unified, highly scalable, modern architecture, Oracle BI EE Plus provides intelligence and analytics from data spanning enterprise sources and applications—empowering the largest communities with complete and relevant insight. Oracle BI EE Plus also bundles key Oracle Hyperion reporting products for integrated reporting with Oracle Hyperion financial applications.

Oracle Business Intelligence Interactive Dashboards.

- This 100 percent thin client solution provides fully interactive collections of analytic content with a rich variety of visualizations. Running on a Web-based architecture, the dashboards provide users with information filtered and personalized for their identity, function, or role based on predefined security rules. The rich, interactive user interface makes the presentation of data intuitive, relevant, and easy to understand. In addition, guided navigation and alerts drive the business user to greater insight and action. As a result, users are guided to make informed and effective decisions that enhance the performance of the entire organization.

Oracle Business Intelligence Answers.

- Providing business users with true ad hoc query and analysis capability, Oracle Business Intelligence Answers can create new analyses from scratch or modify existing analyses in dashboard pages. To free business users from data structure complexity, the metadata layer of Oracle Business Intelligence Answers offers a logical view of metrics, hierarchies, and calculations expressed in understandable concepts. Business users do not need to understand physical data storage to combine data from multiple enterprise information sources intuitively, quickly, and easily.

Oracle Business Intelligence Publisher.

- This integrated pixel-perfect reporting component (formerly known as Oracle XML Publisher) allows the creation of highly formatted templates, reports, and documents such as checks, government forms, and more. It is the most efficient, most scalable reporting solution available for complex and distributed environments. It provides a central architecture for generating and delivering information to employees, customers, and business partners—securely and in the right format. Best of all, users can work with familiar tools such as Microsoft Word or Adobe Acrobat for report layouts.

Oracle Business Intelligence Delivers.

- This powerful, near-real-time, multi-step alert engine can trigger workflows based on business events and notify stakeholders via their preferred medium and channel. This means field sales representatives can receive a short message service alert on their cell phone, warehouse managers get a PDF attachment via e-mail, and financial analysts obtain the report as a Microsoft Excel spreadsheet saved to their shared corporate file system.

Oracle Business Intelligence Plug-in for Microsoft Office.

- Oracle Business Intelligence Plug-in for Microsoft Office integrates business intelligence information from Oracle Business Intelligence Server and Oracle Reports with the Office environment, embedding up-to-the-minute corporate data in Microsoft PowerPoint, Word, and Excel documents. Business users can then share these documents with others for collaborative decision-making. Users authenticate themselves to access Oracle Business Intelligence Server content and, having done so, can access the Oracle Business Intelligence Server metadata to build analyses from objects that can span multiple datasources in the enterprise. This ensures everyone in the organization is working from the single source of truth and not a local copy of the data in Excel.

Oracle Disconnected Analytics.

- The full capabilities and intuitive interfaces of Oracle Business Intelligence Interactive Dashboards and Oracle Answers are also available offline via Oracle Disconnected Analytics. Oracle Disconnected Analytics provides full analytical functionality for the mobile professional, enabling interactive dashboards and ad hoc analysis to be done on a laptop computer while disconnected from the corporate network. Disconnected and remote users can work with their personalized data and maintain all role-based security. Intelligent synchronization of relevant data, analytic metadata, dashboards, and saved selections occurs once a user reconnects to the enterprise network

Hyperion Web Analysis.

- Oracle's Hyperion Web Analysis delivers out-of-the-box OLAP analysis, presentation, and reporting. Hyperion Web Analysis takes advantage of the many powerful and advanced analytics features provided in Hyperion Essbase or SAP BW. The intuitive, easy-to-use interface provides ad hoc analysis capabilities that give Oracle's Hyperion Web Analysis business users the independence to create their own views and rapidly explore a large quantity of data.

Hyperion Financial Reporting.

- Oracle's Hyperion Financial Reporting – System 9 turns enterprise data into operational insights by creating book quality financial reports in a variety of formats for a range of audiences. Compliant with all accounting regulations and external requirements, these easy-to-generate reports can be delivered immediately or scheduled for later distribution to enhance your organization's financial planning, management, and control. With minimal involvement from your IT staff, end users in the finance department can disseminate vital financial information. Oracle's Hyperion Financial Reporting – System 9 supports Oracle's Hyperion Financial Performance Management Applications, Oracle's Hyperion Essbase – System 9, and SAP BW.

Hyperion Interactive Reporting.

- Oracle's Hyperion Interactive Reporting pulls together data from operational or analytic sources to create easily comprehensible charts, pivots, and reports. Oracle's Hyperion Interactive Reporting adds the option of direct access to relational sources without the need to create and access a semantic layer via the Oracle Business Intelligence Server.

Hyperion SQR Production Reporting.

- Oracle's Hyperion SQR Production Reporting – System 9 generates massive reports within specified time frames and connects to a wide variety of enterprise datasources. In the past, combining production planning data residing in enterprise resource planning sources with financial, marketing, and sales intelligence from customer relationship management (CRM) or financial planning sources was next to impossible because most of that data was locked within its own application. With Hyperion SQR Production Reporting, this is no longer the case. By providing access to a wide range of datasources—including leading relational, multidimensional, and legacy databases such as IBM DB/2, SAP R/3, SAP BW, and SQL Server—Oracle's Hyperion SQR Production Reporting – System 9 unlocks that data, providing cross-functional reporting capabilities that let companies surface exceptions and track business process performance.

Oracle Identity Management

- Oracle Identity Management allows enterprises to manage end-to-end lifecycle of user identities across all enterprise resources both within and beyond the firewall. You can now deploy applications faster, apply the most granular protection to enterprise resources, automatically eliminate latent access privileges, and much more. Leverage the Oracle Identity and Access Management Suite in its entirety or deploy individual components of the suite to meet your unique needs.

Oracle Enterprise Single Sign-On

- Enterprise users can enjoy the benefits of single sign-on to all of their applications, whether they are connected to the corporate network, traveling away from the office, roaming between computers or working at a shared workstation. Administrators have the flexibility of easily deploying Oracle Enterprise Single Sign-On into their existing infrastructure without change - no integration or large deployment effort. Oracle

Enterprise Single Sign-On is architected to use any LDAP directory, Active Directory or any SQL database server as its user profile and credential repository. Enterprise Single Sign-On can accept primary authentication directly from the Windows logon, it also works with most industry-leading smart cards, biometrics or token solutions. Confidently delivering stronger password authentication or advanced authentication is simple with Oracle Enterprise Single Sign-On.

Oracle Internet Directory

- Directory services are central to an identity management strategy. Oracle Internet Directory is an LDAP v3 directory that leverages the scalability, high availability and security features of the Oracle Database. Oracle Internet Directory serves as the central user repository for Oracle Identity Management, simplifying user administration in the Oracle environment and providing a standards-based application directory for the heterogeneous enterprise. Additionally, Oracle Directory Synchronization allows Oracle Identity Management to seamlessly integrate with other directories and enterprise user repositories, allowing users to leverage identity information wherever it resides.

Oracle Application Server Portal 11g

- Oracle Application Server Portal 10g Release 2 represents a major step forward in the evolution of Oracle Portal and as such will provide significant value to customers. With this release, customers can enable comprehensive Fusion capabilities for better business agility, unleash the power of portal content management and publishing, enable sophisticated, flexible, and highly performant architectures, and share information instantly with out-of-the-box portal solutions.

Oracle WebCenter

- Oracle WebCenter is an integrated suite of products designed for creating dynamic, highly productive user work environments that take advantage of SOA and enable business users to bring complete context to their daily work tasks. Oracle WebCenter combines the standards-based, declarative development of JavaServer Faces, the flexibility and power of portlets and runtime customization, and a set of integrated Web 2.0 services into a new generation of context-centric, composite applications. Using Oracle WebCenter capabilities, you can build applications that allow users to interact with instant messengers, voice over IP, and other collaborative services directly within the application.

Oracle Developer Suite 11g

- Oracle Developer Suite is the most complete and integrated development environment, combining the power of application development and business intelligence tools in a single suite that is based on the latest industry standards. Oracle Developer Suite lets developers quickly build high-quality transactional applications, which can be deployed to multiple channels - including portals, Web services, and wireless devices, and which can be extended with business intelligence capabilities - including ad hoc query and analysis, high-quality Web reporting, and advanced analysis. Oracle Developer Suite includes JDeveloper, Forms, Reports, Designer, Discoverer, and Business Intelligence Beans.

- Oracle Beehive

  Oracle Beehive is the only unified collaboration solution built for the enterprise. Only Oracle Beehive provides a complete range of integrated collaboration services including conferencing, instant messaging, email, calendar, and team workspaces that can be deployed either on premise or through Oracle On Demand. Its cross-platform, open standards-based architecture supports familiar clients like Microsoft Outlook and can coexist with legacy servers like Microsoft Exchange. Oracle Beehive was designed for the enterprise from its inception to provide the highest level of scalability, manageability, and security while reducing the cost of deployment and administration.

- Oracle Application Express (APEX)

  Oracle APEX is part of the oracle 11g database that allows the creation of application that completely reside in the database.