Leveraging the Oracle/Hyperion Suite of BI: DW Tools to Support an EDW and Human Analytic Solution

Brian Hill
Patrick McCollum

Hitachi Consulting

Background

One of America's largest natural gas utilities recognized the need to accelerate their ability to timely and consistently analyze company performance. The company leveraged the power of a Business Intelligence Competency equipped with Oracle/Hyperion suite of products. OWB is used as the ETL, Oracle 10g the EDW platform and Hyperion 9 BI+ suite the tool to support enterprise analytics. The BI journey began with using these tools to build the new EDW & reporting/analytic solutions to support HR metrics.

Business Background

The business was heavily dependent upon the IT organization to deliver business reporting, generally most often with hand-coded Oracle Discoverer reports. The process of creating reports often took substantial time to create, typically morphing into data extracts that the user community would cobble together to create more meaningful reports. Hence, substantial time was spent gathering data rather than analyzing data.

A great deal of inefficiency was born out due to multiple overlapping reports supporting similar audiences. These ever increasing requests only fostered IT servicing redundant and inefficient requests. This was compounded by the fact that any number of source systems and reporting methods were being utilized to fulfill the needs. Additionally, interest was growing for management dashboards to provide senior management with at-a-glance insight into key company metrics. Financial reporting was the only user self-service reporting available and was provided through Hyperion Essbase, which the users were very pleased with.

The business realized their problems were increasing and decided to get help. Hitachi Consulting was engaged to perform a BI Assessment & Roadmap to provide the client a plan for addressing their needs. The key objectives of the Assessment & Roadmap were to provide an analysis of the overall environment and capabilities with respect to BI ; understand and document high-level reporting & analytics requirements; ensure the involvement and participation of the appropriate mix of business and technology stakeholders; identify the technology necessary to support enterprise reporting and analytics to produce management information, scorecards & dashboards; and evaluate and recommend suitability of existing and available toolsets to meet the BI roadmap requirements.

A Business Intelligence Competency Center (BICC) and the Journey to BI Excellence

A Business Intelligence Competency Center (BICC) program was launched to help the business effectively manage and deliver against BI needs. Some of the benefits include: creates a highly efficient, self-sufficient BI / Data Warehouse delivery engine; establishes enterprise wide leverage for delivering information to decision makers; builds the enabling infrastructure and makes information delivery a core competency by delivering a multi-year roadmap, effective architecture and technology platforms, methodology and data governance.

Some of the guiding principles that governed this program were: integrate data that is easy to access and understand; provide a common tool set to foster a seamless reporting experience across data subject areas and across business groups; deliver a tool set that is intuitive and allows for business self-service through dashboard and scorecard presentation of key metrics; create an organization, methodology and governance processes to support on-going corporate BI growth and maturity; and align BI technology that complements existing IT strategy, skills and size. The principles enable these benefits: boost executive visibility to key corporate performance information; decision-making across the enterprise based on consistent, quality, timely information; analysts move from spending 80% of their time collecting and preparing data to 80% of their time actually doing analysis and making/recommending decisions; empowered end users perform their own ad-hoc reporting and analysis using the common tools; and existing technology and skills investments are cost-effectively leveraged.

An initial focus of the BICC was to execute upon the developed BI roadmap. The initial foundation implementation spotlighted Human Resource Analytics, specifically compensation. It was chosen since HR was a long neglected area of the business and it was a low risk area to start with new BI architecture and tools. Compensation analysis was required to provide key analytics for future strategic metrics requiring Labor Cost Analysis as a function of Operations management. Most importantly, the solution required all major architectural components that were defined in the future vision of the BI roadmap. This encompassed a centralized data in an Enterprise Data Warehouse, with a dependent data mart for HR Compensation analytics. This involved the use of ETL for data integration, relational reporting for Management and Executive dashboard of key metrics requiring robust security measures to protect sensitive data.

Some of the objectives were to supply standard and consistent data for reporting and analysis; enable standard and ad-hoc reporting and analysis for Compensation & Benefits, provide automated dashboard of key HR performance metrics; satisfy the reporting needs of Payroll related to the data requirements for HR on a per pay period basis; support historical and trend analysis and provide processes and tools that ensure the consistent creation, use and reporting of information across departments.

The organization assembled a BICC team that consisted of foundational roles essential with building out a BI solution. These included a BI leader to provide leadership and connection to the BI sponsors, teamed with BI and EDW architects who engineered a best practice BI/EDW architecture and disciplines. The first order of business was to establish the BI tool stack. Once that was determined a delivery team consisting of a data architect, ETL and BI developers were assembled to deliver the first business subject area. As the project was being delivered, best practices and methodologies were created and matured to ensure future BI solution deliverables were produced with a high level of quality and in a consistent and timely fashion.

How did Hyperion 9 Suite changed the way HR and Executive Analysts viewed their information

The ‘front-end’ tool selection focused on three short listed vendors Oracle OBIEE, Hyperion System 9 (Note: This took place before the Oracle acquisition of Hyperion) and Business Objects XI. Several demonstrations and Q&A sessions were held to vet out the right tool. Hyperion System 9 was the chosen solution since it provided the best opportunity for integrated reporting between existing Essbase and future relational reporting needs. Other factors included: cost (software & hardware), vendor demos, and core functionality vs. user requirements.

We know the business was impacted immediately as evidenced by the following testimonials: "The Hyperion Interactive Reporting Tool has put users more in control of the data they need to access and utilize on a daily basis. The volumes of data that previously needed to be extracted from our Payroll system and provided to us by IT can now be accessed solely through the use of straightforward queries within Interactive Reporting. Response time to provide data in support of internal and external data requests has been significantly reduced; in some cases from days to hours. The Hyperion Interactive Reporting tool has made an immediate and noticeable impact in our department.” If was not enough, here is another one: “…Interactive Reporting provides a robust capability which allows our users to be in control of the data they need… It drastically reduces the time resource and man power needed to provide our service of delivering reports and analysis. The dashboard functionality enables a much more dynamic graphic interface of trends in our workforce which in return, helps management keep a pulse on various

aspects of our employee population. The Interactive Reporting tool has certainly made an immediate and noticeable impact in our organization and consequently, moves us to a more advance level of our technology spectrum."

The use of the tool enabled greater insight into the HR aspects of the business making it possible to alter benefits and compensation. It further empowered the user with the ability to view data in an analytic fashion versus typical static report view. Clearer understanding of the business terms by way of a semantic layer reduced confusion about data lineage and content. IT assistance was reduced since the user was more self-sufficient. Finally information was accessed in seconds versus days. The use of System 9 lessened the dependence on Discoverer reports, provided a usable and friendly interface which enabled the user to create interactive reports and dashboards. System 9 IR was used to develop dashboards and reports to support Health & Welfare, Retirement, Compensation and Payroll analysis.

Even though there was significant success, behind the scenes there were some challenges specific to configuration and development of System 9 components. Most notable were substantial manual intervention in the configuration process – wizard processes were challenging; BI+ services failed on occasion; IR Studio desktop application dashboard fatal error would often corrupt BQY file (backup BQY frequently); Drill-up/drill-down in Workspace HTML version of IR dashboard didn't work, but worked in Web Client. Hyperion is fixing in version 9.5l; workspace scheduler stops running without reason and Workspace errors when changing file properties.

From a user's perspective, the experience was positive. It was easy for the end user to view the semantic layer to support the build of their reports. They were able to easily create tables, charts and pivot tables. Unquestionable benefits were they became increasing independent from IT assistance, realized substantially increased access to data and a generally positive attitude about the tool experience.

A solid benefit of the product is the use of 'business friendly metadata'. IR Studio by default presents the physical model of the underlying data store. The team exploited the capability of IR Studio to present business friendly metadata, which was maintained in an external metadata repository. The data management team employed effective data governance practices to keep metadata up-to-date. The metadata connection was configured at the OCE level.

In general, the steps taken were rather straightforward, yet required a rigor to ensure that the metadata kept in the model was kept in synchronization with the BI tool. The first step is to create metadata specific tables in the Oracle data mart. Each of those tables contained column definitions and typical column characteristics. Next the data names, definitions and attribute characteristics were exported from CA All-Fusion (Erwin) into metadata structures within the data mart. Ultimately, the metadata was synchronized with System 9 repository.

The second step is to import the metadata structure and data into IR Studio, through the database connection wizard. Filter the metadata tables that will be part of the OCE connection. The next step is to configure the metadata connection essentially linking in the metadata tables to the metadata aspect of the tool. Next, link the metadata column definition and table remarks. Finally the metadata presentation is laid out so that the user will be able to view and use.

The 'Master Data Model' concept was used to leverage reusable queries rather than creating the same or similar query over and over. To promote increased conformance and ease of maintenance, BQY content was organized in a subject-oriented manner. Only a limited number of master data models were created. Also, metatopics were created as part of the master data model to promote ease of query building. Master Data Model and metatopics were kept as generic as possible to promote wider usage (i.e. do not insert Filters, etc. that may not be applicable to all queries).

Finally, keep in the mind that there are some good habits to follow when creating reports and dashboards. Some general suggestions are: pay particular attention to the "Query Processing Order"; do not use (or minimize usage of) "Custom SQL" and "Computed Columns" as this does not lend itself well to ease of centralized maintenance; minimize custom coding in Interactive Reporting Reports – included computed columns. Where possible, retrieve pre-computed values from the source (e.g. Data mart) – this has dual benefit: lower reporting authoring and maintenance efforts and improved conformance of metrics.

What were the success & challenges of using Oracle Warehouse Builder

The ETL tool selection focused on 3 short listed vendors. Oracle Warehouse Builder (OWB) 10gR2; Informatica and IBM DataStage. Why OWB? The existing ETL base was hand-coded PL/SQL stored procedures and was very challenging to maintain and enhance. After reviewing other ETL tools, OWB was chosen primarily due to various factors. It was an ETL tool, hence it help the team avoid pitfalls of hand-coded programs; the technology platform was an all Oracle environment; OWB had improved features from prior versions; transformation features of OWB was sufficient for the job and finally it contained Oracle ERP connectors.

Benefits to the business, albeit not readily seen by the user, is by aligning to the one vendor enabled the firm to realize cost savings and simplified support. It boosted the team's capability to quickly and efficiently add subject data to the Data Warehouse and Data Marts. Transparent and consistent data transformation enabled a higher level of quality in the solution. Streamlined and efficient ETL processing lessened burden on infrastructure all leading to lower development and maintenance costs.

How is OWB being used? First, it is used to extract data from the Oracle ERP (HR modules), with minimal transformation, ultimately loading HR data into a staging area. Next, the staged data is transported, leveraging OWB capabilities, to transform and normalize data into the Data Warehouse (Oracle). Third, it sourced data from the Data Warehouse, organizing and aggregating into dimensional or star schema form (Oracle). In tandem with AppWorx, OWB orchestrated process flows to load data into the various BI structures

Approximately 210 OWB maps and 34 process flows were created to load the BI environment. Most loads processed in just a few minutes, with an exception of monthly processing of larger volumes of data, which took 45 minutes. In general the performance was good despite the very complex transformation processes. Notable is the fact the same daily, weekly and monthly processes were used to preload the data warehouse with 12+ years of data, taking about 2+ days to complete. Combined with solid ETL practices and OWB capabilities, the implementation went relatively smooth.

As with any tool, there are noteworthy observations or lessons learned. OWB is no exception. The initial stage in installing and configuring Oracle Warehouse Builder was fairly cumbersome due to the various versions, compatibility, patches and "bug-fixes". OWB has its own learning curve. The documentation was lacking and it was a challenge to find answers. Work a rounds were hard to find (lots of google'ing and message boards). Our project development cycle was extended due to working the issues and bugs in the product. There are really few OWB experts.

Specific functional notables are: it can not create reusable or shareable objects; column deletion from mappings disorders input and output groups especially in the EXPRESSION Operator; cannot use LONG data in OWB; columns typically need to be changed to Varchar2 after importing tables from source system; sequence object generates unused numbers since it generates sequences for all incoming rows; and ordering objects and conditions in the process flow is critical to proper execution. Continuing on: a common process flow issue is invalid status for process flow activities; since the code is generated, it is easy (maybe too easy) to "customize" the code to suit ones requirements. Finally, backing up the repository is very critical.

When in the tool, there are some best practices that should be considered. When importing metadata objects, always match by 'Names' rather than the universal identifier. Do not purge the deployment data as this has been known to corrupt the OWB repository. Before you modify or delete any object, run impact or lineage analysis. Impact displays dependent objects and mappings and lineage shows what objects were used to populate. Always synchronize the target table 'from and to' table operator by object name not by position or ID. Frequently save the mappings during design time. Many times the mapping will corrupt with an unknown JAVA error. Backup the development project at least once a day using MDL exports keeping several versions. When deleting a column from operators in a mapping - first delete the data flow connectors; next delete the field and finally reconnect. Always arrange the table Target Load Order in the following order: TYPE1 INSERT, TYPE1 UPDATE, TYPE2 INSERT and TYPE2 UPDATE.

For better job schedule flow and restart ability, create process flows in an organized approach. Create master process flows, gradually decomposing to sub-processes through each layer (Stage, EDW and DM). When integrating with an external scheduler (process flows prefixed by PF), leverage OWB schedule for most of the work (typically those prefixed with SPF). Transition order does matter in Process flows. Join error condition transitions first and then success. Error condition is in the '0' order and success is '1' order.

Additional good practices to follow: Give each developer their own OWB username (i.e. don't all log on as REP_USER). Apply patches, as they become available. Define a security policy (FGAC vs. simple security on repository objects). When moving a project from one system to another, always import the locations first; make appropriate connection changes and import one location object at a time. Purge the execution audit details and optimize the repository whenever the repository browser is very slow. Run the repository database in ARCHIVELOG mode to allow for possible point-in-time recovery. Capture issues in the log files by piping the OWB outputs to a file.

There are some pretty good practices to following when using any ETL tool. Create a set of naming standards at the outset for the mappings, tables, operators, process flows or other objects. Build and automate scripts to audit or reconcile row counts and key business values between Source, EDW and Data Marts. Log all the ETL job runs in a Control table. Maintain business rules within an ETL rule table rather than hard coding business values. Create maps that will manage 'dates by source' to allow for the filtering of source data by selective time frames. Create modular, single process mappings for each target table. Design and build each mapping with performance in mind. Do not hide functionality in your Mappings. Avoid using views and PL/SQL procedures in your maps. Use Filter Operators to restrict data sets, not predicates in Join Operators. Create map templates to ensure consistent creation of maps.

Conclusion:

The combination of people, the appropriate tools, a sound architecture and the implementation of sound practices led to a successful deployment of the BI foundation and initial business focus area. A well planned roadmap was the blueprint that made it possible for the business to orchestrate a successful launch to their BI program. The deliberate execution of the roadmap focused the selection of tools to those that were absolutely appropriate for the environment, best leveraging existing investments and skills to support the long term health of the BI initiatives. In the end, the testimonials say it all.