Teaching Note for Continental Airlines Real-time BI Case

"Continental Airlines Flies High with Real-time Business Intelligence" is a comprehensive example of how an organization used real-time data warehousing and business intelligence applications in its turn around from a poorly performing to an industry leading organization. The case is not heavy on technical details, but the essence of the technology is explained. Thus, the case can be used for a variety of courses and in different degree programs from data warehousing to DSS/BI and from undergrad to executive. Discussion of the case can be led by using the following nn questions and exercises. Included with each question or exercise are some of the points that we believe are most important to raise in the discussion.

1. Describe "active" data warehousing as it is applied at Continental Airlines. Does Continental apply active or real-time warehousing differently than this concept is normally described?
   - Continental uses "right-time" instead of real-time, meaning they have an appropriate mixture of data feeding and reporting cycles based on what refresh rates the business needs, how timely reported results need to be, and what is the most cost effective
   - Continental employs many real-time feeds of data into the data warehouse, and many examples of real-time, on-demand (prewritten and ad hoc) reporting routines and querying tools (e.g., Teradata QueryMan, and Hyperion Brio)
   - So far, Continental has limited instances of direct links from analytics against the data warehouse back to operational systems (e.g., results from analytics on the data warehouse affecting the interactions with customers on the Continental on-line sites); thus, Continental has many examples of real-time, but few examples of active data warehousing

2. In what ways does real-time data warehousing fit with the Continental strategy and plans?
   - To change the company, Continental decided to be on the leading edge of applying high-value technologies and business concepts, and real-time warehousing distinguishes Continental among their competitors
   - Each element of the Go Forward Plan can be translated into consequences for data management and business intelligence
   - The Go Forward Plan was hampered in general by the fact that information (from a single version of the truth) was not available within Continental to make critical decisions, yet they were data rich with many sophisticated, separate operational and decision making systems
   - The initial application of data warehousing concentrated on one of the most critical areas for Continental to go from worst to first—revenue management; thus, from the beginning, data warehousing was positioned as supporting business strategy, which gave the initiative considerable visibility
   - The initial application also addressed one of the most troublesome operational problems facing Continental—the inability to understand market and customer behavior, including a customer's itinerary from origin to destination (what has led to the PNR database)
   - Data warehousing benefited from being associated with the initial wins for the business turn around effort, thus giving it a foothold for playing a role in future investments in implementing the new business strategy

3. Describe the benefits of real-time data warehousing at Continental.
   - Each application, from the initial ones to the most current, have shown cost savings or revenue enhancements (e.g., the initial 'demand-driven dispatch" application is attributed with a $5M dollar revenue increment)
   - Many of the revenue benefits have focused on gaining more revenue from high-value (high-profit) customers via higher-quality service, which leads to greater loyalty, and eventually more business with greater revenue
   - Other benefits have reduced fraud from travel agents and customers
   - Inside and outside of Continental, benefits have also been in greater passenger and industry security
   - There has also been savings in data center management due to the efficient processes utilized and the consolidation of decision support services possible from the enterprise data warehouse
   - In general, the benefits come from providing real-time, actionable information to support tactical decision making and business processes
   - In total, Continental has achieved an ROI of over 1,000 ($30 million investment for over $500 million of increased revenues and cost savings)

4. What did the data warehousing group do right that has led to the successful deployment of (real-time) data warehousing and BI within Continental?
   - Developed a warehouse architecture that could grow and scale to meet these new real-time and operational needs
   - Developed prototypes or demonstration applications to show potential end users, to get them excited about data warehousing, and to give them ideas about new applications that could be built (part of an R&D effort to prove the concepts)
   - As stated in a previous discussion question, they linked applications to business strategy initiatives and insured there was business benefits for each application
   - The data warehouse operates consistent with organizational culture; the warehouse provides employees with information and tools that they can use to do a better job, so they can "go out and change the world."

5. What elements of the data warehousing environment at Continental are necessary to support the extensive end-user BI application development that occurs?
   - Data exist in the data warehouse from sources that are trusted by end users
   - Users have access to and are trained in tools to access and manipulate data
   - Help from data warehousing staff is readily available and friendly
   - Metadata is kept current and is easily accessible by end users via the web
   - Graphics are used, when appropriate, for data display, making it easier for users to understand and interpret the complicated data being presented
   - In many cases, users can run "what-if" scenarios to determine the impact of decisions

6. Why does Continental believe that a 3rd normal form enterprise data model is important?
   - Lesson Learned #8 addresses this question
   - The comprehensive nature of an enterprise data model naturally results in a 3NF data model

- A 3NF model is easier to evolve (as new subject areas are added) without reorganization and changes
- A 3NF data model is easier to administer because it eliminates duplicated data, and hence avoids potential inconsistencies in meaning and value

7. What special issues about data warehouse management (e.g., data capture and loading for the data warehouse (ETL processes) and query workload balancing) does this case suggest occur for real-time data warehousing? How has Continental addressed these issues?
   - Date and time management is amplified because of the finer granularity of data; although the techniques available to handle date and time stamping are the same for any data warehouse, users in a real-time environment have to be better trained to deal with time in queries and the large volume of data requires qualifications on time be handled as quickly as possible (e.g., use of active flag quickly identifies the most recent values)
   - Customized views (with indexes, joins, and aggregations) significantly improve query performance and reduce the load on the data warehouse
   - With the extensive number of on-line, real-time users, views also provided an extra level of security against access to unauthorized data
   - Data loads come in via many different routes and methods (push-pull, subscribe, queues, trickle-batch, etc.), so generalizable components to handle data loading are used to save the effort of starting from scratch to develop each new loading process
   - The large volume of constant data loading means that it is not humanly possible to watch all ETL processes, so automated watchdog applications are used to alert data warehouse staff via pagers when their attention is needed for some anomaly
   - Data for loading are put into standardized queues, from which pre-written load utilities pull data for loading into the data warehouse, no matter what the source of the data are
   - Because there are data loads, tactical queries (prewritten and ad hoc) and strategic queries, each with different patterns of data warehouse use, specific priorities are given to the different types of loads against the warehouse. Priorities also change by type of day. Higher priority is given to queries that require the fewest data warehouse processing resources.

8. What generalizable lessons learned can be gleaned from this case? Are there other lessons for you beyond those listed in the Lessons Learned section?
   - Ten specific lessons are outlines in the Lessons Learned section; these lessons can be applied to the development of real-time data warehousing in any organization
   - Another theme of the case suggests an additional lesson, which is that data warehousing, which can be viewed as a discretionary investment, will be better received if it is linked with business strategy