

Real-time Business Intelligence: Best Practices at Continental Airlines¹

Written by Hugh J. Watson (University of Georgia), Barbara H. Wixom (University of Virginia), Jeffrey A. Hoffer (University of Dayton), Ron Anderson-Lehman (Continental Airlines), and Anne Marie Reynolds(Continental Airlines)

Data management for decision support has moved through three generations, with the latest being real-time data warehousing. This latest generation is significant because of its potential for affecting tactical decision making and business processes. Continental Airlines is a leader in real-time business intelligence and much can be learned from how they have implemented it.

The movement to real-time is the latest development in business intelligence (BI) and data warehousing. Real-time data warehousing provides the data that is required to implement real-time BI. By moving to real-time, firms can use BI to affect current decision making and business processes. This capability is especially important for customer-facing applications, such as those found in call centers and check-in processes, and helps firms become more customer-centric. Terms such as the “real-time enterprise” and the “zero latency organization” are often used to describe firms that use real-time BI.

The purpose of real-time BI is to increase revenues and decrease costs. Companies that successfully implement real-time BI can dramatically improve their profitability. For example, Continental Airlines, which is discussed later, has taken a \$30M investment in hardware, software, and personnel to generate over \$500M in revenue enhancements and cost savings, resulting in a ROI of over 1,000 percent.

To be successful with real-time BI, organizations must overcome both organizational and technical challenges. On the organizational side, there must be executive sponsorship and support, initial and on-going financial support, governance processes put in place, BI and data warehousing personnel with the requisite skills, changes in business processes and acceptance of use of real-time data by organizational personnel. On the technical side, new hardware and software must be acquired and implemented, processes and procedures for supporting and managing real-time data feeds from source systems must be established, data must be quickly transformed and loaded into the warehouse, and the data must be analyzed and made available to operational systems and personnel.

In this article, we provide frameworks and discuss key issues that are helpful in understanding real-time BI. We then use Continental Airlines as a case study of highly successful real-time BI. In 2004, Continental won The Data Warehousing Institute’s prestigious Best Practices and Leadership Awards. We briefly discuss Continental and their business strategy that led to implementing real-time BI, describe an application that illustrates Continental’s use of real-time BI, and focus on the technical issues associated with Continental’s implementation of real-time BI.

¹ We gratefully acknowledge Teradata, a division of NCR, who provided funding for this research.

Putting Real-time BI in Perspective

Before we explore real-time BI at Continental, it is useful to put real-time BI and data warehousing in context. In particular, real-time data warehousing is the latest of three generations of data management for decision support. It is also useful to sort out what real-time really means as it applies to BI. And finally, we discuss latency, its impact on the value of data, and how it requires both technical and organizational solutions.

Three Generations of Data Management in Decision Support

The use of data for decision support can be conceptualized as moving through three generations. The *first generation* emerged with decision support systems (DSS) in the early 1970s. It was recognized that DSS applications required a repository of data, some of which was sourced from operational systems, but also other data, such as external data. The data was customized for the specific DSS that was developed. This was a very application-centric approach, with the data supporting a single or a few related applications. It did, however, show the critical role of data in decision support. In his seminal work on decision support, Sprague provided the Data-Dialog-Models (DDM) paradigm, which recognizes data as one of the cornerstones of DSS.¹

The *second generation* emerged in the late 1980s. Firms in the telecommunications, retailing, and financial services industries built data warehouses to store vast amounts of customer and sales-related data. Companies in these industries remain leaders in terms of the size of their warehouses and how the warehouses are used. Unlike DSS in the first generation, data warehouses tend to be data-centric. While a single or a few applications may be used to help make the business case for the warehouse, the data is modeled to support a variety of applications. The term “single version of the truth” is commonly used to describe the official repository of data that applications are supposed to use.²

In 2000, the *third generation* began with the movement to real-time data warehousing. The major reason this development is significant and worthy of being a new generation is the changes in the way that warehouse data is used. Previously, the data was primarily employed to understand what had already happened and to predict what would happen in the future. Its use for influencing real-time decisions and current operations was limited. With real-time data, current decisions and critical business processes, such as customer-facing and supply chain applications, can be significantly enhanced.

Real or Right Time?

For many people, the “real-time” term is synonymous with “instantaneous.” This interpretation, however, is incorrect when applied to data warehousing. While some warehouse data may be captured and entered into the warehouse in seconds or minutes, much of it is not. For example, some source systems, such as a legacy COBOL program that is updated once a month, can never be more real-time than when last updated. Some data may be prohibitively expensive or difficult to make real-time. Most importantly, there may not be a business need for real-time data. Data only needs to be as fresh as the business requirements. For these reasons, some people prefer the

“right time” term.³ We use them simultaneously and recognize that real-time does not always mean instantaneous.

The Latency and Value of Data

In most cases, the value of data decreases rapidly as it ages. Stating it differently, low latency (i.e., fresh) data has more value than high latency data. This is why the movement to real-time BI is appealing.

Richard Hackathorn provides a useful perspective on latency as applied to data warehousing.⁴ He identifies three kinds of latency; see Figure 1. *Data latency* is the length of time between when an event occurs and when the associated data is stored in the data warehouse. *Analysis latency* is the time between when the data is stored and when it is analyzed and made available to applications and users. *Decision latency* is the time from when the information is available until some action is taken on it. These three sources of latency are additive and result in total latency.

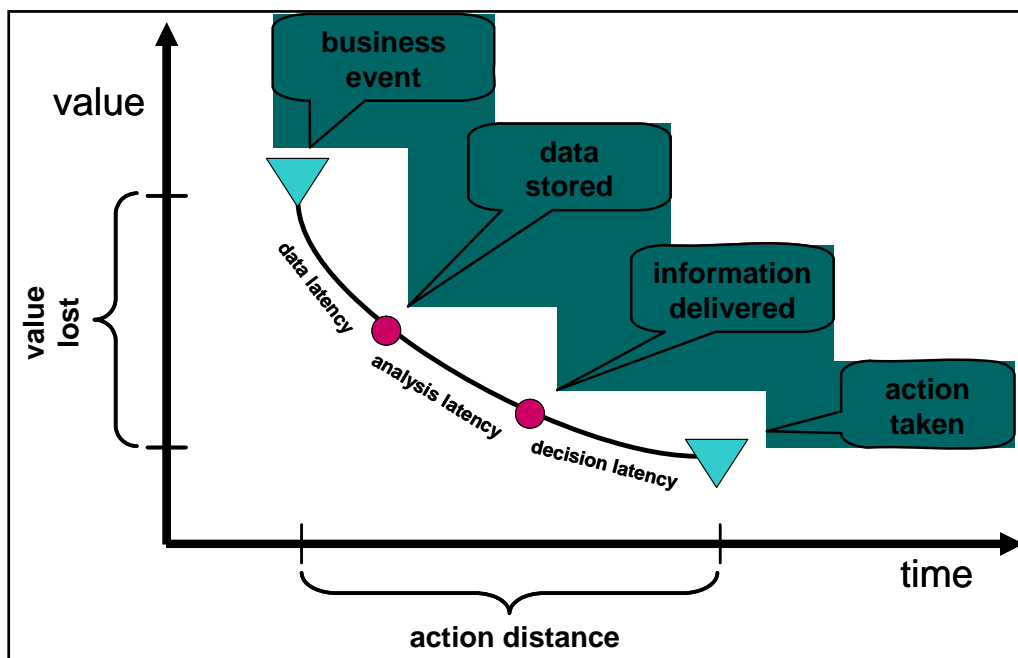


Figure 1: Types of Latency⁵

Reducing data and analysis latency depends primarily on technical solutions. Recent developments in real-time data warehousing provide help in this regard. However, reducing decision latency requires changes in business processes and how people use information in performing their jobs. Providing fresher data does not create business value unless it is used in a timely manner. Dealing with decision latency is usually more challenging than data and analysis latency.

Continental Airlines

Continental Airlines is a leader in real-time BI. It has received numerous awards for its work, including in 2004, The Data Warehousing Institute's Best Practices and Leadership Awards.⁶ Continental's experiences with real-time BI illustrate the challenges, solutions, and business value associated with real-time BI.

About Continental Airlines

Continental Airlines was founded in 1934 with a single-engine Lockheed aircraft on dusty runways in the American Southwest. Over the years, Continental has grown and successfully weathered the storms associated with the highly volatile, competitive airline industry. With headquarters in Houston, Texas, Continental is currently the USA's fifth largest airline and the seventh largest in the world. It carries approximately 50 million passengers a year to five continents (North and South America, Europe, Asia, and Australia), with over 2,300 daily departures, to more than 227 destinations. Continental, along with Continental Express and Continental Connection, now serves more destinations than any other airline in the world.

Continental's Business Strategy

Continental was in trouble eleven years ago when Gordon Bethune became CEO.⁷ There were ten major US airlines, and Continental consistently ranked tenth in the Department of Transportation metrics used to monitor the industry's performance: on-time arrivals, baggage handling, customer complaints, and denied boardings because of overbooking. Not surprisingly, with this kind of service, Continental was in financial trouble.

Bethune and Greg Brenneman, who was a Continental consultant at the time, conceived and sold the Board of Directors on the Go Forward Plan. It had four interrelated parts that had to be executed simultaneously.

- *Fly to Win.* Continental needed to better understand what products customers wanted and were willing to pay for.
- *Fund the Future.* It needed to change its costs and cash flow so that the airline could continue to operate.
- *Make Reliability a Reality.* It had to be an airline that got its customers to their destinations safely, on-time, and with their luggage.
- *Working Together.* Continental needed to create a culture where people wanted to come to work.

Under Bethune's leadership, the Go Forward Plan, along with a re-energized workforce, has helped Continental make rapid strides. Within two years, it moved from "worst to first" in many airline performance metrics, including on-time performance, lost baggage claims, and customer satisfaction.

After this success, Gordon Bethune and his management team raised the bar with a new vision. Instead of merely performing best, they wanted Continental to be their customers' favorite airline. The First to Favorite strategy built on Continental's operational success and focused on treating customers extremely well, especially the high-value customers.

The Role of Information Technology

The movement from “worst to first” was only partially supported by information technology. Historically, Continental had outsourced its operational systems to EDS. These included mainframe systems that provided a limited set of scheduled reports and no support for ad hoc queries. The airline lacked the corporate data infrastructure that a broad range of employees could use for quick access to key insights about the business.

In 1998, the decision was made to develop an enterprise data warehouse that all employees could use for quick access to key information about the business and its customers. The CIO at the time, Janet Wejman, recognized that the warehouse was a strategic project and brought the development and the subsequent maintenance and support in-house. She believed that the warehouse was core to Continental’s business strategy and should not be outsourced.

The data warehouse provided a variety of early, big “wins” for the business. The initial applications for pricing and revenue management were followed by the integration of customer information, finance, flight information, and security. They created significant financial lift in all areas of the Go Forward Plan.

However, when Continental moved ahead with the First to Favorite strategy, it became increasingly important for the warehouse to provide real-time, actionable information to support tactical decision making and business processes. Fortunately, the warehouse team had anticipated and prepared for the ultimate move to real-time. Real-time meant that the warehouse team had to introduce real-time feeds of data into the warehouse, extract data that the warehouse produced and incorporate them back into legacy systems, and open the warehouse to tactical queries with sub-second response time requirements. In preparation, the team had developed a warehouse architecture that could grow and scale to meet these new real-time and operational needs. While not all applications required real-time data, many did. In 2001, real-time data became available in the warehouse.

Real-time BI Applications

Continental’s real-time applications fall into the following categories:

- Revenue management and accounting
- Customer relationship management
- Crew operations and payroll
- Security and fraud
- Flight operations

The objective of revenue management is to maximize revenue given a finite set of resources. An airline seat is a perishable good, and an unfilled seat has no value once a plane takes off. The revenue accounting area seeks to quickly and accurately record the revenues that Continental generates, including estimating the revenues from any flight as soon after “the wheels are up.”

The Marketing group employs customer relationship management (CRM) in order to increase revenues, profits, and customer service by knowing customers exceptionally well (e.g., customer value, flying preferences) and giving them great service. Continental's Marketing group uses the warehouse for customer segmentation and target marketing, loyalty/retention management, customer acquisition, channel optimization, and campaign management.

The Crew Operations group is concerned with issues related to pilots and flight attendants. It is involved in crew pay, crew scheduling, crew performance, and crew efficiency. The data warehouse is used in conjunction with all of these activities at varying levels.

Continental uses its warehouse to identify reservations that are not in fare and contract compliance and to profile suspicious booking and ticketing transactions. It is also used to support airline security efforts.

Real-time applications have been developed for flight operations. The objective is to get people to their destinations safely, on-time, efficiently, and with their luggage. This is where customers have either a good or bad flying experience, and Continental works hard to provide consistently excellent service. One of the flight operations applications is the Flight Management Dashboard, and it is described next as an in-depth example of how Continental uses real-time BI.

Flight Management Dashboard

The Flight Management Dashboard is an innovative set of interactive graphical displays developed by the data warehouse group, with input from the operations staff at Continental's hub in Newark, New Jersey. These displays are intended to help the operations staff quickly identify issues in the Continental flight network and then manage flights in ways to improve customer satisfaction and airline profitability.

Some of the dashboard's displays help Operations to better serve Continental's high-value customers. For example, one of the displays is a graphical depiction of a concourse, which is used to assess where Continental's high-value customers are or will be in a particular airport hub; see Figure 2. The display shows gates where these customers have potential gate connection problems so that gate agents, baggage supervisors, and other operations managers can assess where ground transportation assistance and other services are needed so these customers and their luggage avoid missing flights. In Figure 2, it can be seen that Flight 678 is arriving 21 minutes late to Gate C37. Three high-value customers need assistance in making their connections at Gate C24 and will have only 12 minutes to make their flight (see the box at the bottom right-hand side). Five high-value customers need to get to Gate C29 and will have 20 minutes.

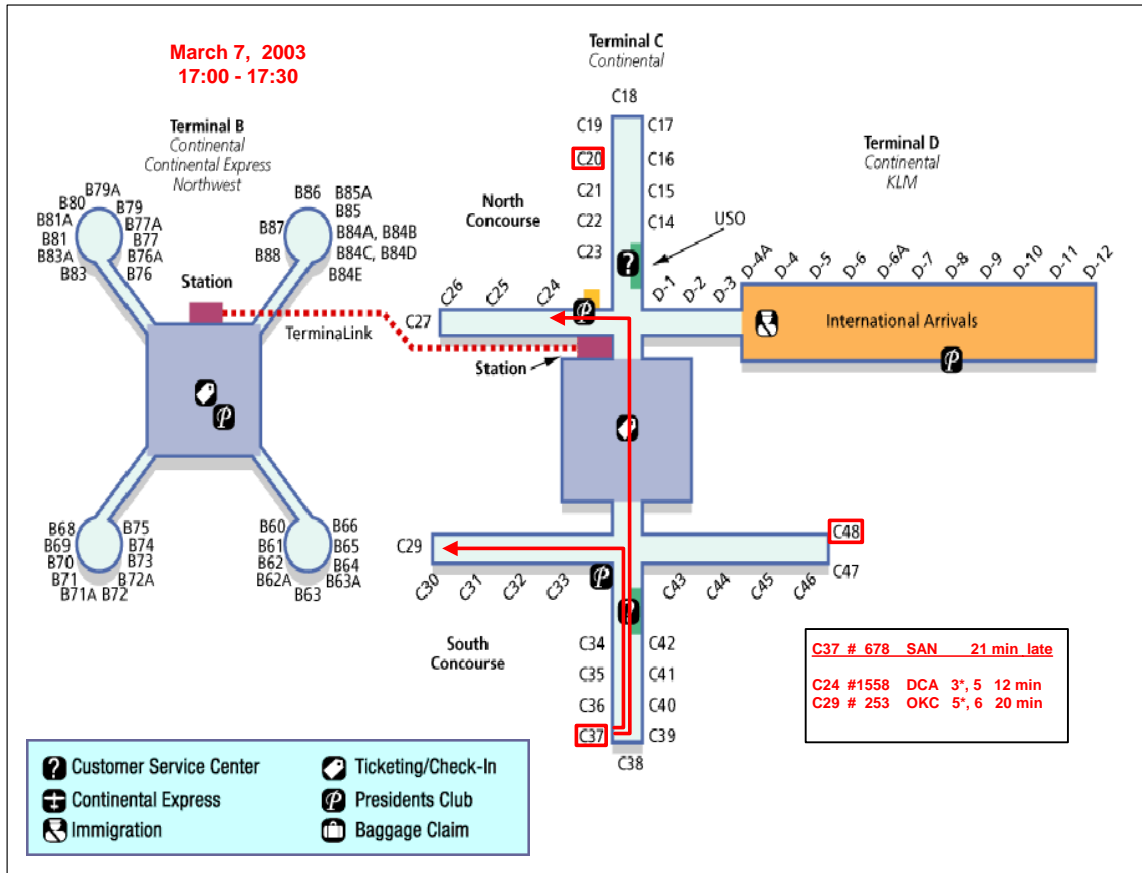


Figure 2: Concourse Display of High-Value Customer Activity

On-time arrival is an important operational measurement at Continental. Therefore, another critical set of dashboard displays helps Operations keep the arrivals and departures of flights on time. One display shows the traffic volume between the three Continental hub stations and the rest of their network; see Figure 3. The line thickness between nodes is used to indicate relative flight volumes and the number of late flights so that the operations staff can anticipate where services need to be expedited. The ratio of the number of late flights to the total number of flights between the hubs is also shown. The operations staff can click on the lines and drill down to see individual flight information. Another line graph summarizes flight lateness. Users can drill down to more detailed pie charts that show degrees of lateness, and then, within each pie, to the individual flights in that category. Another chart concentrates on flights between the US and Europe and the Caribbean, and can show similar critical flight statistics. In all of these elements of the dashboard, high-level views can be broken down to show the details on customers or flights that compose different statistics or categories.

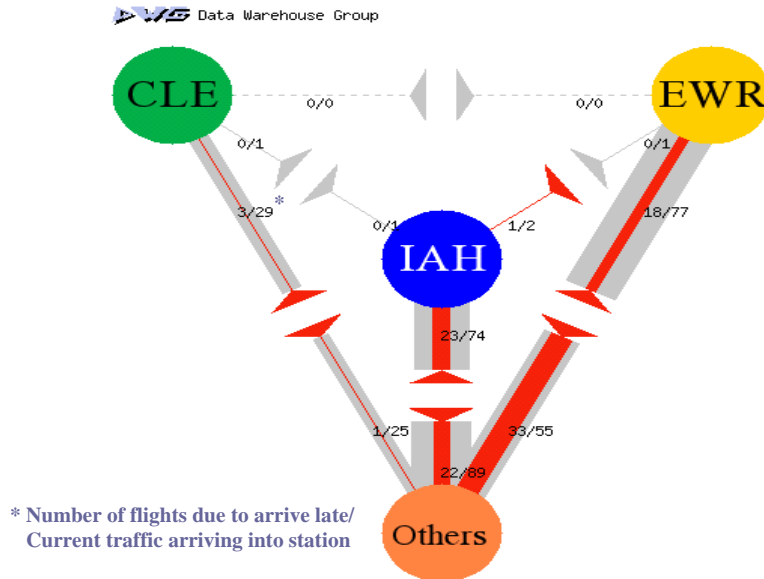


Figure 3: Display of Flight Lateness From/To Hubs

The Architecture for Real-Time Business Intelligence

The architecture underlying the Flight Management Dashboard is not trivial and is built on a real-time data warehousing foundation. At the core is an 8-terabyte enterprise data warehouse running on a 3 GHz, 10-node Teradata 5380 machine. The warehouse supports 1,292 users across Continental who access 42 subject areas, 35 data marts, and 29 applications. Sixteen of these users are “power users” of the Flight Management Dashboard. More users access the Flight Management Dashboard infrequently for special decision support needs.

The basic architecture of the warehouse is shown in Figure 4. Data from 25 internal operational systems (e.g., the reservations system) and two external data sources (e.g., standard airport codes) are loaded into the warehouse. Some data (e.g., customer value analysis) is fed from the warehouse back into the operational systems. All of the data are stored at the lowest level of detail in the Teradata database.

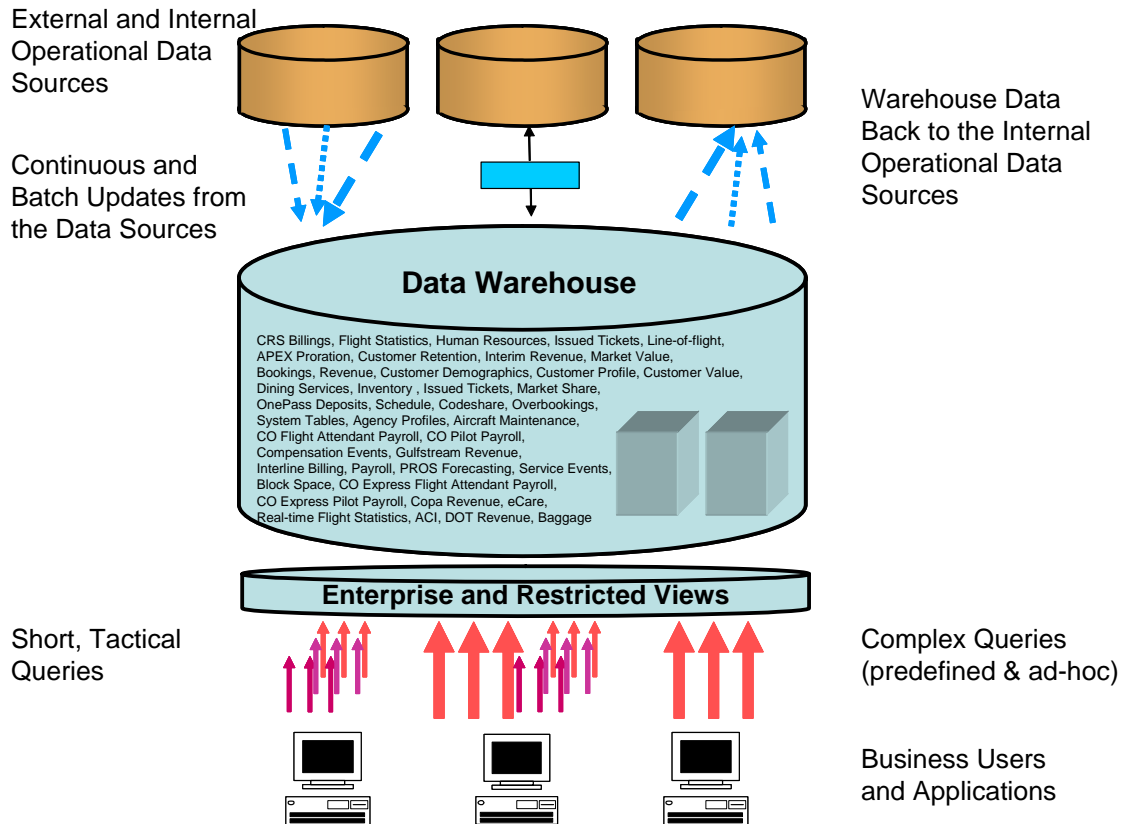


Figure 4: The Data Warehouse Architecture

Real-time Data Sources

In order to deliver the Flight Management Dashboard application with the functionality described earlier, the warehouse team needs data about a flight’s current status combined with passenger information. This requires two primary real-time data sources – the satellite feeds that are transmitted from airplanes and the central customer database. Because these data sources exist in two very different technical environments, they must be extracted from the sources in different ways.

The flight data (called FSIR, or flight system information record) is sent real-time from the airplanes via satellite to an operations control center system, which supports the command center for Continental where the actual flights are coordinated. The data warehouse team took advantage of the command center infrastructure already in place, and set up a Windows NT machine that “listens” to the FSIR data and feeds the appropriate data (as they occur real-time) into a warehouse queue.

Whereas the warehouse team uses a process to “pull” the flight data, it instead uses a “push” approach with the passenger data. The passenger data sits in central Oracle customer database that is accessed and updated by Continental’s reservations system, the One Pass frequent flier

program, Continental.com, and customer service applications. Every time a change is made to a customer record in the customer database, an Oracle trigger is activated, which sends the update as XML-encoded data to a queue for loading into the warehouse.

Although the data sources have clear technical differences and need different extraction processes, Continental created an infrastructure called the Service Bureau that allows the various sources to be captured and monitored using a single, reusable infrastructure and then loaded in a uniform way. The data capture, monitoring, and loading processes are described in the next two sections.

Data Capture and Monitoring

Continental's vision from the start was for the data warehouse to be highly automated. To achieve this, the warehouse team built a set of shared services that run on 10 Windows-based servers called the Continental Service Bureau. The Service Bureau was built using object-oriented design techniques, and it automatically loads and manages the warehouse. Its object orientation facilitates efficient, parallel, scalable, and re-startable transformation processes. The Service Bureau contains a number of components, which are listed in Figure 5.

The Service Bureau consists of:

- a set of components and libraries that implement the common services,
- scheduler, file transfer, and data conversion programs,
- a watchdog service to ensure other services are up,
- a system for paging on-call personnel through the Skytel service,
- an e-mail queuing service that connects to the corporate e-mail system,
- COM + components that provide a uniform Data Access Layer for web and other clients for logging and connection pooling to minimize database connection usage,
- a set of services that transform and load reservations and operational data into the warehouse in real-time,
- a process that monitors the performance of the data warehouse, keeps history, and alerts the warehouse staff to exceptionally long-running queries,
- automatically generated meta data that is refreshed and published on the Intranet,
- an SQL Server operational database that stores configuration and status data,
- control of Teradata load utilities, and monitoring of MVS jobs through 3270 emulators, and
- real-time display of what each service is doing in each transformation server machine.

Figure 5: Components of the Service Bureau

The Service Bureau's scheduler has about 50 tasks that run at intervals from minutes to hours, and others that run at specified times. Many of these tasks check whether a load is ready to run, and if so, start it. Most of the loads depend on either a file becoming available through the Internet or the completion of another load. The dependencies are quite complex.

The Service Bureau manages a series of automated alerts that page the on-call warehouse staff member when some process needs human intervention. If the on-call person fails to respond and fix the problem, another page is sent to a secondary support person.

The Service Bureau monitors the number of items in the various work queues for the real-time loads, and if the queues fill up beyond certain preset thresholds, it issues an alert, paging the on-call staff. The Service Bureau also notifies the on-call staff member when any data warehouse load fails, or a data feed is not available in time, or a batch load process is taking too long and is projected to finish beyond its allotted time. It also monitors the latency of the data in the warehouse and issues an alert if it gets beyond a predefined value.

For batch loads, the Service Bureau notifies the users of the data when their data has been loaded and is ready for the day.

A Standard Data Loading Process

More than half of the daily data added to the warehouse is loaded throughout the day, with varying degrees of latency. This includes the flight and passenger data required for the Flight Management Dashboard application.

Using the Service Bureau, data feeds are treated similarly regardless of their source. As feeds evolve from batch to real-time, there are minimal changes to the loading process. Regardless of latency, data feeds are transformed and sent to queues, which then are loaded using the real-time loading utility provided by Teradata. Figure 6 provides an overview of the data loading process.

Because the standard loading process is generic and reusable, as users demand more and more real-time data, and as real-time feeds become available, incorporating real-time data into the warehouse is relatively simple.

1. Continuous data loads use software, such as CDC (Change Data Capture), or special hooks or triggers on the operational systems to capture events of interest.
2. Special queue middleware, such as MQ series, Oracle AQ, and MSMQ move the data to dedicated transformation servers.
3. Each update to the operational systems “triggers” a record to the queue.
4. Service processes continuously monitor these queues and transform each piece of data that gets loaded into the warehouse using the Teradata continuous load TPump facility.
5. Strategic and tactical queries can run the same time that data is being loaded.

In the case of batch updates, the operational systems create nightly files that are FTP'd to the transformation servers where they create an event that causes the transformation programs to run. These programs transform the data and push the transformed data through the Teradata Fastload and Multiload utilities to load the data into the warehouse.

Figure 6: The Real-time Data Loading Process

Managing Mixed Queries

Traditional warehouses are optimized to support strategic queries that require many table joins and aggregations. However, this is not the case with queries associated with real-time BI. For example, as the Operations staff accesses the Flight Management Dashboard application, they need to know a specific flight's status at a certain point in time (usually the current time), and they expect an immediate response to their query. This type of query is tactical and, therefore, requires a different query response strategy than a strategic query to ensure appropriate performance.

The challenge is to optimize real-time warehouses so that strategic and tactical queries can co-exist. Users of the Flight Management Dashboard application need to "hit" the same data as users of more strategic applications without degrading overall performance of the warehouse. Continental's optimization strategy aligns with their overall vision to "keep things simple," and this approach has worked well.

The warehouse team takes advantage of the data warehouse's existing priority scheduler application to help manage the mixed query loads. They have concrete rules that determine the priority settings for users and applications. The tactical queries that access single records are set to *high* priority. These queries usually come from applications, such as the Flight Management Dashboard that require instantaneous response time. Daytime data loads are set to *low* priority. And, users who perform ad-hoc queries are given *medium* priority access.

The Service Bureau continuously monitors all of the queries to identify those that are using too much of the system's resources. For example, when queries run longer than one hour, the monitor service alerts someone on the data warehouse team. Query responses that are not needed immediately are set to the lowest priority. Warehouse team members work with users to help them understand how to optimize inefficient queries.

The Data Warehouse Team

Continental has 15 people on its data warehouse team. They are responsible for managing the warehouse; developing and maintaining the infrastructure; data modeling; developing and maintaining extraction, transformation, and loading (ETL) processes; and working with the business units. The organization chart for the data warehouse staff is shown in Figure 7.

The Data Warehouse Director reports to the Chief Information Officer. The warehouse staff is located in Miami and Houston. When the warehouse initiative began, Continental filled the Director position with a highly experienced woman from Miami who had previously worked with excellent data warehouse professionals in the area, and she was able to hire them if they could continue to live in Miami. The Houston and Miami groups work as a team and share the infrastructure development and maintenance work, including building the processes that source data for the warehouse. The most technical people on the team have degrees in computer science.

Everyone on the data warehouse team wears many hats, including providing operational support for the warehouse. However, they also have specialized roles. For example, the Master Data Modeler has ultimate control over the warehouse’s enterprise data model.

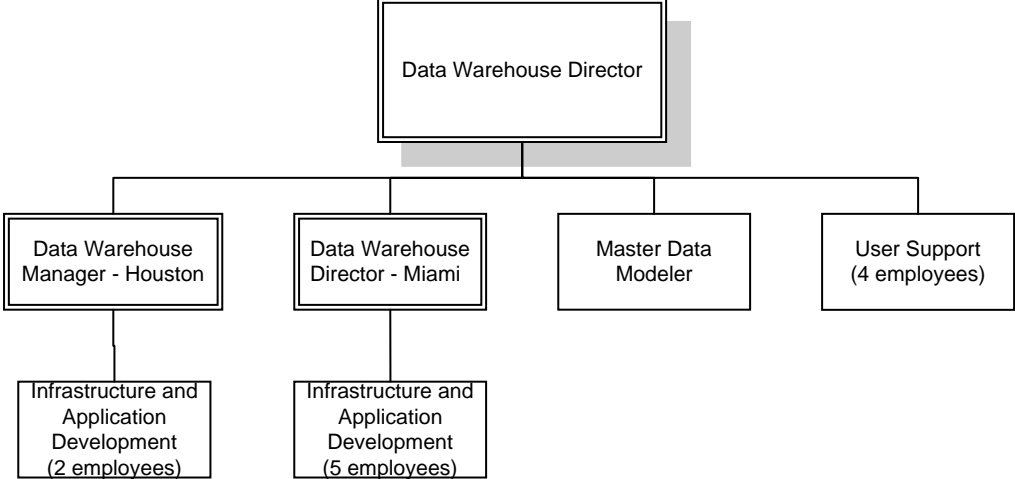


Figure 7: The Data Warehouse Organization Chart

Four team members have specific support roles. Three work with the Revenue Management, Marketing, and Crew Operations areas. The fourth is responsible for supporting end user business intelligence software and miscellaneous tasks. All of the support people originally worked in the user areas they now support and, therefore, are experts on the data for those areas. They assist rather than build applications for the business areas. Several team members have extensive work experience with operational systems, which has helped them in implementing real-time data warehousing.

Considerations with Real-Time Business Intelligence

Over time, business needs will drive organizations to evolve into the third generation of decision support data management. As the movement occurs, there are some nuances to real-time business intelligence that organizations need to consider.

Changes in latency is an evolution.

Business needs increasingly require more real-time data; however, to reduce latency requires technical solutions, a strong business case, and process change. Putting these three factors in place may take time, and may also require an evolutionary approach where organizations move to real-time data in incremental steps. A key to success is for the organization to recognize that latency needs will change and then construct an architecture that can handle the eventual movement to real-time.

Continental recognized the need and planned for real-time BI at the outset of its data warehousing initiative. For example, the Service Bureau infrastructure supports data loads using

queues, and this approach works for both batch loading and real-time feeds. The Service Bureau is set up to continuously monitor all warehouse processes and respond rapidly to issues (e.g., data loading failures, long query response times); the Bureau automates all processes all the time, regardless of how often they run. All of this has helped transition Continental to real-time business intelligence painlessly.

It is important to define what real-time means to your organization.

A flight is considered “late” within the airline industry if it takes off or arrives more than fourteen minutes from its scheduled time. When the data warehouse director discussed “real-time” with the business, she used fourteen minutes as her latency window because that number made sense from a business perspective.

Specifically, the Flight Management Dashboard relies on satellite data that is captured and loaded into the data warehouse every couple of seconds. With service levels of fourteen minutes or (typically) less, the Operations users can feel confident in the data as they make real-time decisions that affect current flight operations.

Users need help initially understanding the potential of real-time BI.

At first, the Operations staff did not understand why they would need a business intelligence application to help manage their day-to-day tasks. The data warehouse team had to think of ideas and present them to the users to stimulate dialog and excitement. The data warehousing staff began by creating the Concourse Map screen, which is a key piece of the Flight Management Dashboard application. When the users saw how data could be depicted in graphical ways (e.g., as an actual concourse with colors and lines that have special meaning), they came up with their own ideas for how real-time data could help them operate the hubs better. Once users appreciate what is possible, they are more likely to say: “Help me change the way that we do business.” At Continental, the current challenge is to find the time to support the ideas that users have.

Real-time BI requires the automation of ETL processes.

The feeding of real-time data to a data warehouse is analogous to the running of a manufacturing plant. The processes should be as automated as possible. There should be minimal human intervention unless monitoring systems detect a condition that requires attention. Furthermore, the processes should be flexible and reusable so that changes can be made easily. Continental’s vision from the start was for the warehouse to be self-maintaining. The warehouse team developed automated processes (in the Service Bureau) that minimize human intervention and can be used across source systems and with new source systems that come on line.

Strategic and tactical queries must be carefully managed to successfully co-exist.

Historically, data warehouses have focused on providing strategic decision support (when needed, operational data stores have supported operational decision making). Users analyze data and then take appropriate actions. With the emergence of real-time technologies, decision

support has expanded to the tactical level. With tactical decision support, the information is used to support operational decisions.

Strategic and tactical decision support systems have different characteristics, yet must co-exist in the same warehouse environment. Strategic decision support typically involves the analysis of large amounts of data that must be “sliced and diced” in various ways. Tactical decision support often requires repeatedly accessing only a limited amount of data. The performance requirements (i.e., response time) are greater for tactical than strategic decision support, since there is a smaller window of time for when the information is useful. The resource demands for strategic decision support are often greater and more varied than for tactical decision support (e.g., data mining versus accessing a customer value score).

The successful co-existence of strategic and tactical decision support requires business and technical solutions. On the business side, priorities must be set for the processing of queries from users and applications. For example, a data-mining query should have a lower priority than a tactical query. There must be capacity planning for each class of query. On the technical side, there must be a query manager that recognizes the priorities of the queries, monitors queries, defers long-running queries for later execution, and dynamically allocates query resources.

Real-time BI blurs the line between decision support and operational systems.

In many companies, the operational and decision support systems are “islands.” Each has its own staff, processes, and standards. The only significant point of interaction is the extraction of data from operational systems. And, the systems are managed in different ways. Operational systems have more stringent service level requirements and more sophisticated business continuity processes in place.

The Flight Management Dashboard is one example of the blurring of the distinction between decision support and operational systems. As the Operations staff become more reliant on the application to manage and make operations decisions, the application will become more “mission critical.” Continental will have to ensure that the application is treated in an operational way; otherwise, business processes will break down.

Focus on changing downstream decision-making and business processes that are enabled by real-time BI.

When the Flight Management Dashboard was first proposed, the person who was responsible for creating the overnight performance reports was not convinced that real-time information was needed. Perhaps more likely was that he was very comfortable with generating reports from the mainframe and was hesitant to give up his area of expertise. It was only after he had worked with the real-time data for a while that he saw its value and integrated it into his work. Now he is actively developing an application that will even better present real-time information.

Clearly, when organizations delve into the real-time BI world, they need to apply change management practices and work with decision makers to ensure that benefits of real-time BI are realized.

Conclusion

As some organizations identify the need to move to the third generation of decision support, they may see the migration as “crossing a chasm” because of the dramatic changes to technology architecture and organizational processes that will need to occur. Those organizations that recognize the progression of decision support early on, may instead be able to lay groundwork to facilitate a comfortable transition to real-time BI over time.

As we see with Continental, there are clear technical, business, and process changes that must be put in place to enable real-time BI. The Continental case offers an approach for moving to the third general of decision support from which other organizations can learn.

First, find a clear business need for decisions that requires real-time data. At Continental, Operations could directly impact costs and revenues by altering gates, rerouting baggage, and making other operating modifications – but only if the data reflected the business at the current time. It was obvious that only real-time BI would be able to enable decision making for the Operations staff.

Next, invest in an architecture that can scale, and then automate as much of the environment as possible. Continental benefited from a mature data warehouse that had been architected to receive data feeds from a system of queues. And, they had created the Service Bureau to manage and monitor the technical processes automatically. This foundation did not need to change as the latency of data feeds moved from batch to real-time.

Finally, perform a value assessment to support the investment in real-time BI. Continental regularly assesses how applications are impacting the company’s bottom line, and for this reason, they do not plan to feed all data warehouse applications with real-time data. The cost and effort are prohibitive for some applications when compared to the benefits. However, those applications that can leverage real-time BI by impacting business processes to create value to the organizations will represent the third generation of decision support.

Notes

¹ R.H. Sprague, Jr., “A Framework for the Development of Decision Support Systems,” *MIS Quarterly*, 4, no. 4 (December 1980), pp. 1-26.

² P. Gray and H.J. Watson, *Decision Support in the Data Warehouse*, Upper Saddle River, NJ, Prentice-Hall, 1998.

³ C. White, “Now is the Time for Real-Time BI,” *DM Review*, September 2004, pp. 47-54.

⁴ Richard Hackathorn, “The BI Watch: Real-Time to Real Value,” *DM Review*, 14, no. 1 (January 2004).

⁵ Richard Hackathorn, “The BI Watch: Real-Time to Real Value,” *DM Review*, 14, no. 1 (January 2004).

⁶ Additional information about the real-time BI best practices at Continental is described in R. Anderson-Lehman, H.J. Watson, B.H. Wixom, and J.A. Hoffer, “Continental Airlines Flies High with Real-Time Business Intelligence,” *MIS Quarterly Executive*, 3, no. 4 (December 2004), pp. 163-176.

⁷ The story of Continental’s turnaround can be found in G. Brenneman, “Right Away and All at Once: How We Saved Continental,” *Harvard Business Review*, 76, no. 5 (September/October 1998), pp. 162-74.