

White Paper

## Going Real-Time for Data Warehousing and Operational BI

### Enabling Real-Time Data Integration

**Abstract:** Gone are the days when Data Warehouses were just for reporting, analytics, and forecasting. Today, more companies are working hard to make their data warehouses operational and active -- and thus more critical to the business. An important requirement is to enable lowest possible latency in which new data is delivered to the data warehouse, ideally in real time. And this new requirement for “fresher data” must often be given careful consideration when investments in batch-oriented ETL technologies already exist.

Why is fresher, real-time data so crucial? In order to determine and influence what should happen next, the enterprise data warehouse needs to know what is happening right now. There are many data integration technologies that serve the data acquisition needs of a data warehouse, and the demand for low-latency data is causing the IT group to evaluate a range of approaches: intra-day ETL batches, mini-batches, EAI, ELT, and real-time change data capture (CDC) techniques. The challenge is to determine what solution or combination of solutions will meet the trend for faster data delivery needs, which will propel the move to operational and active data warehousing.

This paper will review the trends that are driving the data warehouse towards real-time data requirements, examples of the business value, evaluation of data acquisition technologies and approaches, an overview of GoldenGate's technology solution for real-time data warehousing (and GoldenGate's flexibility to integrate with existing ETL implementations), and real-world examples of GoldenGate customers who have been leading the trend in this area.

## Table of Contents

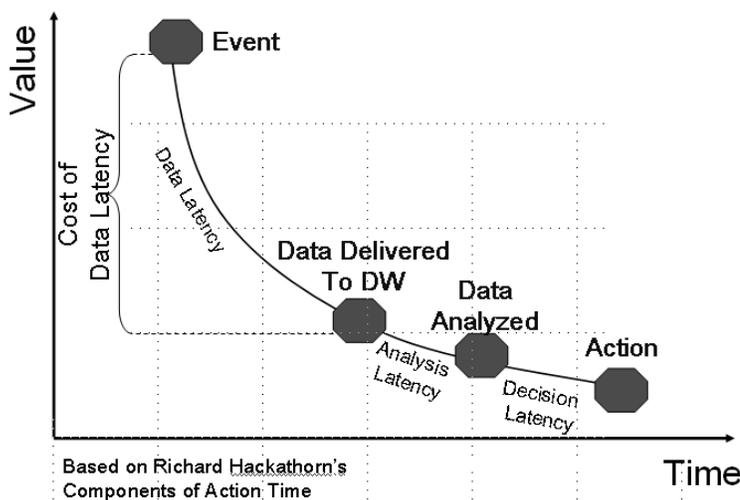
1. **Introduction: The Data Warehouse and “Real Time”**
2. **The Data Warehouse Evolution: From Reporting to Active Decision Support for Front-Line Users**
  - a. Strategic Stages: Reporting, Analytical, Predictive
  - b. Tactical Stages: Operational, Active
3. **Data Acquisition Approaches for the Real-Time Enterprise**
  - a. Scripts
  - b. ETL
  - c. EAI
  - d. Change Data Capture (CDC)
  - e. Data Transformations – Where Do They Belong? (ETL or ELT?)
4. **GoldenGate TDM for Real-Time Data Warehousing**
  - a. How GoldenGate TDM Works
  - b. Combining GoldenGate TDM with ETL
5. **GoldenGate Customers: Real-Time Data Warehouses in the Real World**
6. **Summary & Contact Information**

## 1. INTRODUCTION: THE DATA WAREHOUSE AND “REAL TIME”

Business time is increasingly moving towards real time – in order to grow their competitive advantage, organizations should seek to capture and respond to business events faster and more rigorously than ever. Today, the preponderance of competitive advantage comes from the effective use of information technology. Therefore from an IT standpoint, the key tool of the trade for enabling business intelligence is the enterprise data warehouse, coupled with an enterprise analytics framework.

Across the enterprise, each facet of the business gathers data through an assortment of business activities and ultimately delivers it to a central data repository. This repository is the data warehouse – where data is captured, aggregated, analyzed, and leveraged to drive better decisions. The quality of these decisions depends not only on the sophistication level of the analytics applications but also on the underlying data. Data has to be accurate, relevant, and complete. But most importantly, it has to be **timely**. Timely data ensures better-informed decisions.

The lifecycle of a data record through enterprise analytics starts with a business event taking place. Data acquisition technologies deliver the event record to the data warehouse. Analytical processing helps turn the data into information, and a business decision leads to a corresponding action. To approach real time, the duration between the event and its consequent action needs to be minimized. Typically, it is the data acquisition process that introduces majority of the latency.



Industry analysts have recently reported on this trend for business intelligence and data warehousing – largely driven by the leaders in various industries who have deployed real-time capabilities and demonstrated the clear value.

At the Gartner 2007 Business Intelligence Conference, analyst Mark Beyer presented that “by 2010, more than 90% of all production data warehouses will have intra-day updates or better, and more than 50% will be continuous (0.7 probability).”

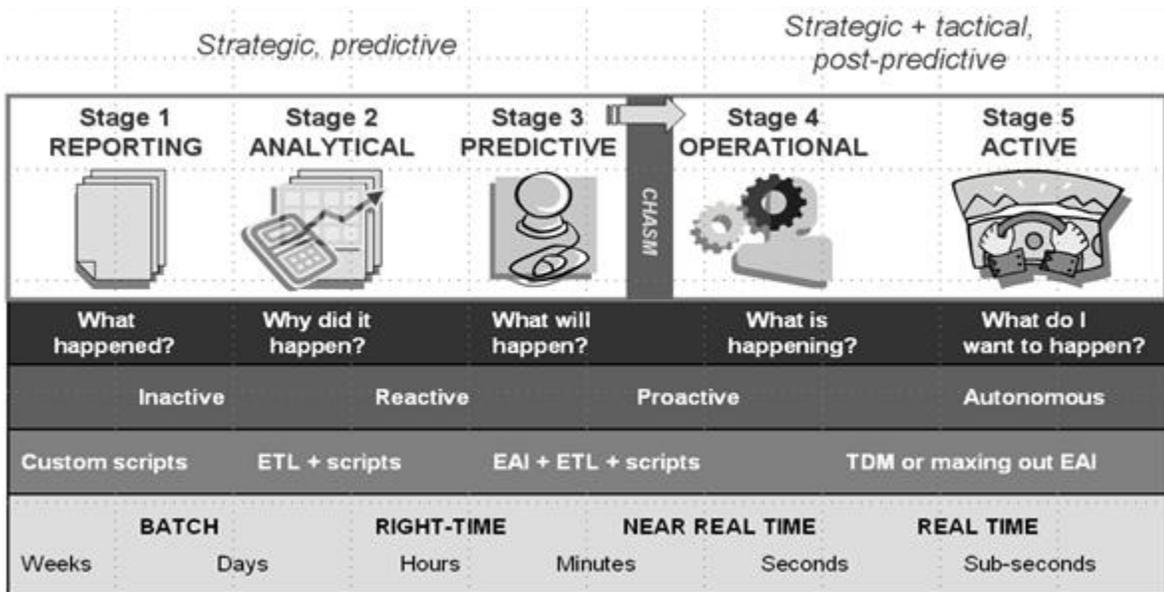
Similarly, Forrester released a ForresterWave™ for Enterprise ETL, Q2 2007 which stated that “while ETL technology is still heavily used for data warehousing (DW) and business intelligence (BI) initiatives, information and knowledge management professionals demand additional data integration capabilities from their ETL vendors to support increasingly complex data integration challenges. **These challenges may often include real-time or near real-time data integration requirements.**”

So where does a business go from here? If the trend is for more “timely” data, there is a great deal of subjectivity around how fresh the data needs to be for the business needs. Moving from nightly batch loads of data to a real-time data approach requires not just business justification, but technology solution availability and effectiveness.

This paper will address both the business reasons to move to real-time data warehousing, as well as describe some of the common data integration approaches – with emphasis on Transactional Data Management (TDM) technology and case study examples of its benefits to companies who have deployed it for BI solutions.

## 2. THE DATA WAREHOUSE EVOLUTION: FROM REPORTING TO ACTIVE DECISION SUPPORT FOR FRONT-LINE USERS

Traditional data warehouses have been strategic-only resources that help create reports, analyze events, and predict what might happen in the future. Today’s data warehouses are not only strategic but also tactical; adding mission-critical decision support to their workload.



The data warehouse evolution started 15 years ago when Bill Inmon defined it as "a subject-oriented, integrated, time-variant and non-volatile collection of data in support of management's decision making process." Businesses wanted to know "what happened". During its early stages, data warehousing was all about reporting historic information aggregated from transaction processing systems with the goal of providing a unified, integrated view of business activities. Data was predominantly brought to the warehouse using custom scripts. Batch oriented, these scripts were executed monthly, weekly, and sometimes daily to update the warehouse.

**Stage 1: Reporting:** *Company ABC determines how many sweaters were sold in December across their stores in the US and by knowing that, they could make better educated purchasing and stock management decisions for the upcoming year.*

As data warehouses evolved, businesses were no longer satisfied with knowing what had happened – they wanted to understand the underlying reasons. The questions grew more sophisticated:

**Stage 2: Analytical:** *Beyond just knowing how many sweaters were sold in December, Company ABC wants to understand how many sweaters were sold in New England during a snow storm.*

This second stage of data warehousing witnessed the birth of data integration technologies such as Extract, Transform and Load, or the now familiar, ETL. Much larger volumes of data had to be brought from a wide variety of sources, transformed, correlated and loaded to the data warehouse. Often a combination of

custom scripts and ETL tools executed data integration tasks weekly or daily. Users could request occasional ad-hoc queries and obtain segmented results so that they knew how many sweaters were sold in Boston following a snow storm and separately, cross reference that data with how many sweaters were sold in their other geographical locations.

By the turn of the millennium, data warehouses grew to hold unprecedented volumes of data. Having a larger amount of data meant it was now statistically meaningful for the warehouse to make predictive analysis.

**Stage 3: Predictive:** *If there was a snow storm in New England, Company ABC could predict that sweater sales would jump by 40% in Boston but remain unchanged in Los Angeles.*

The time-value of data started to gain importance; rather than monthly and weekly updates, data warehouses started receiving daily and hourly updates. In addition to custom scripts and ETL, technologies such as EAI (enterprise application integration) started to augment data integration processes to ensure timely acquisition of data. Data quality, aggregation, profiling, and metadata management became integral to data warehousing. Business intelligence technologies entered the mainstream and business users became increasingly proactive with the aid of their budding enterprise analytics frameworks.

As we are halfway through the 21<sup>st</sup> century, many data warehouses have become “post-predictive”. This is the stage that leads to what Wayne Eckerson defines as the “chasm” (DM Review, November 2004) where the data warehouse starts moving from monitoring business processes to driving the business and ultimately the marketplace. The fundamental mismatch between transactional systems and analytical systems has led to different approaches to close the gap, which paved the path to burgeoning data marts, operational data stores, and ultimately operational and active data warehouses.

To fulfill the true potential of the enterprise data warehouse and “cross the chasm”, the data warehouse has to become a mission-critical enterprise resource.

**In the “post-predictive” stages, the data warehouse becomes operational. The business needs to know what is happening right now in order to determine and influence what should happen next** and start leveraging enterprise data as the organizational memory. The enterprise data warehouse can now support not just strategic business intelligence, but fuel dynamic, well-informed decision-making for front-line operational personnel. Operational data warehouses enable business users to become **proactive**.

**Stage 4: Operational:** *With the aid of an operational data warehouse, a business analyst with Company ABC can track how many sweaters are being sold as it starts to snow in Boston -- and call the store manager to move merchandise from the warehouse to the shelves, AND quickly launch an offer for discounted gloves at the checkout to drive up-sell revenue.*

At this stage, data acquisition in the warehouse has to approach as close to real-time as possible – there is a lot less value in knowing about and reacting to the snow storm 24 hours after the fact.

In the final stage, the data warehouse goes active. As real-time data feeds the data warehouse and matches pre-defined business patterns, business actions are automatically triggered. The active data warehouse auto-initiates actions to systems based on rules and context to support business processes.

**Stage 5: Active:** *At this stage, when the SKU numbers for certain sweaters start to sell at a rate that exceeds the average by 40%, the active data warehouse automatically notifies the store to move additional merchandise to the store shelves and start displaying discounted gloves at a particular price – with little human intervention.*

With active data warehousing, not only does the data need to be delivered in real time, but it also has to maintain its transactional context. Assume that a corporate customer plans ahead and purchases 50 sweaters in August as a December holiday gift, the data warehouse should be able to recognize this as an anomaly – a feat that requires transactional context of the business event to be kept in tact. Maintaining the integrity of data transactions is a data acquisition concern for active data warehousing.

Providers of Enterprise Data Warehousing technologies, such as Teradata, describe other examples of the value of operational and active data warehousing:

- > Users numbering in the hundreds to tens of thousands would benefit: Gate agents, cashiers, dock

workers, bank tellers, salespersons, customer service/call center agents, customers, and suppliers

- > Intelligence for operational execution would include product look ups, individualized customer offers, transaction exceptions, supply chain visibility, event detection and notifications
- > The EDW would help personnel make excellent “small decisions” that add up to better overall competitive advantage that match overall business strategy:

“Should I expedite this package based on the criteria I see here?”

“Should I extend a special offer to this customer at this moment?”

“Should I make changes to our current marketing campaign based on this morning’s results?”

Another consideration for Stage 4 & 5 enterprise data warehouses is the **availability of the warehouse** for supporting 24x7 tactical decision making. Once real-time data empowers the operational warehouse, any planned or unplanned downtime now directly affects the business processes it supports. As Gartner’s analyst Mark Beyer also recently stated, “fewer than 15% of data warehouses in 2007 have been designed to provide high availability, failover, disaster recovery and the remaining components of mission-critical systems” (Report: *Operational Analytics and the Emerging Mission-Critical Data Warehouse*, May 14, 2007).

### 3. DATA ACQUISITION APPROACHES FOR THE REAL-TIME ENTERPRISE

There are numerous technologies that serve data acquisition needs, but only a few offer real-time data delivery with low impact and without relying on batch windows.

“Right-time” – another popular phrase – should always be a business decision, not a technology constraint. Unless the technology is real-time, the business cannot choose “real-time” as the “right-time.” Technology should deliver real-time capabilities and let the business user choose the most appropriate time as the “right time.” Right-time should be a component of decision latency; a user preference determining when the action should take place.

The selection criteria should also focus on important considerations such as data quantity, frequency of data movement, acceptable latency, data integrity, transformation requirements, processing overhead, and outage windows/business continuity.

Traditional data acquisition approaches likely include scripting, ETL (extract, transform, load), EAI (enterprise application integration), and Change Data Capture (CDC). Scripts and ETL are batch oriented in data delivery, while EAI and CDC support continuous data capture.

Attribute	Scripts	ETL	EAI	CDC
<b>Data Volume</b>	Medium	Very High	Low	High
<b>Frequency</b>	Intermittent	Intermittent	Continuous	Intermittent or Continuous
<b>Latency</b>	Medium to High	Medium to High	Low	Low
<b>Transactional</b>	No	No	No	Varies by Offering
<b>Transformations</b>	Intermediate	Advanced	Basic	Basic
<b>Processing Overhead</b>	Intermittently High	Intermittently High	Continuous & Medium	Continuous & Low
<b>Batch Outages</b>	Yes	Yes	No	No

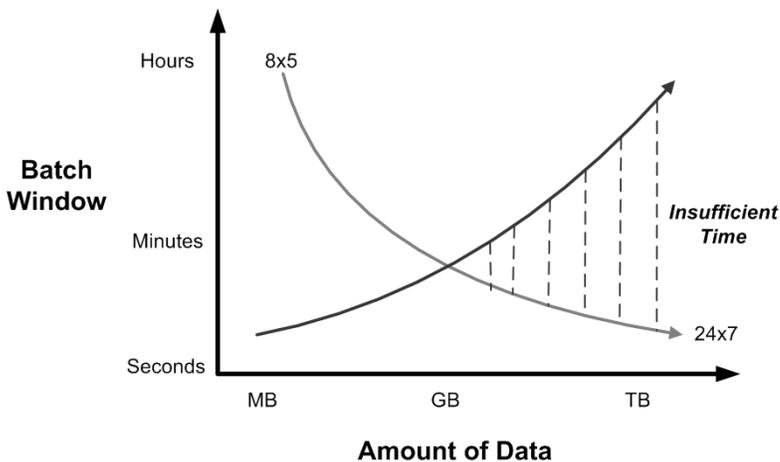
### Scripts

Scripts are a quick solution to data integration. However, they pose many challenges such a drain on developer resource time and effort, in addition to administrative challenges such as manageability, documentation, and SLA compliance. On the other hand, they are flexible and economical to develop and modify. Almost every operating system can invoke scripts from their built-in scheduling facilities.

### ETL

ETL is the ideal solution for the loading of very large volumes of data. ETL also offers advanced transformation capabilities. As for data acquisition, ETL tasks are executed intermittently – typically during nightly maintenance windows when the data sources are acquiesced, to ensure that data sources don't change during data acquisition and lead to inconsistencies across OLTP systems and the data warehouse. ETL may also take advantage of change data capture (CDC) technologies for real-time data acquisition from the source systems, but that data is still staged on the ETL server and eventually batch loaded into the target warehouse.

However, batch windows are disappearing. Significant growth in transactional data in OLTP systems, combined with the need to keep those systems and the operational data warehouse highly available at or near 24x7, leaves little to no time to take the batch window. At some point, a different data integration approach must be considered. Running several batch loads per day (mini-batches or intra-day batches), in an attempt to increase the frequency of refreshing the enterprise data warehouse, means outage and overhead trade-offs while never truly achieving “real time.”



### EAI

Originally designed and intended for application integration, EAI solutions have evolved into a real-time data integration solution. EAI solutions continuously deliver data between source and target systems, provide guaranteed data delivery, feature advanced workflow support, and facilitate basic data transformations. However, EAI imposes data volume constraints; since the original intent is integrating applications rather than integrating data, EAI is designed to invoke applications and move instructions and messages. Nevertheless, with its ability to move data in real-time and maintain the integrity of the data through the integration process, EAI provides real-time data acquisition capabilities required for some types of operational and active data warehousing needs.

### **Change Data Capture (CDC)**

CDC technologies identify and capture changes made to enterprise data sources, and then deliver those changes to target systems. CDC does not require downtime or batch windows for ETL, provides more timely information to the data warehouse, and can improve performance on those source systems. However, not all CDC solutions are created equal, not all are transaction-aware, and those different offerings mean different implications on the overall solution's latency, scalability, flexibility, and data integrity and recoverability. While CDC eliminates the batch window and heavy overhead on the source, some CDC technologies still operate end-to-end in batch mode with a pull approach -- meaning the ETL tool periodically requests to receive a batch of all new changes made since the last request, then proceeds from there.

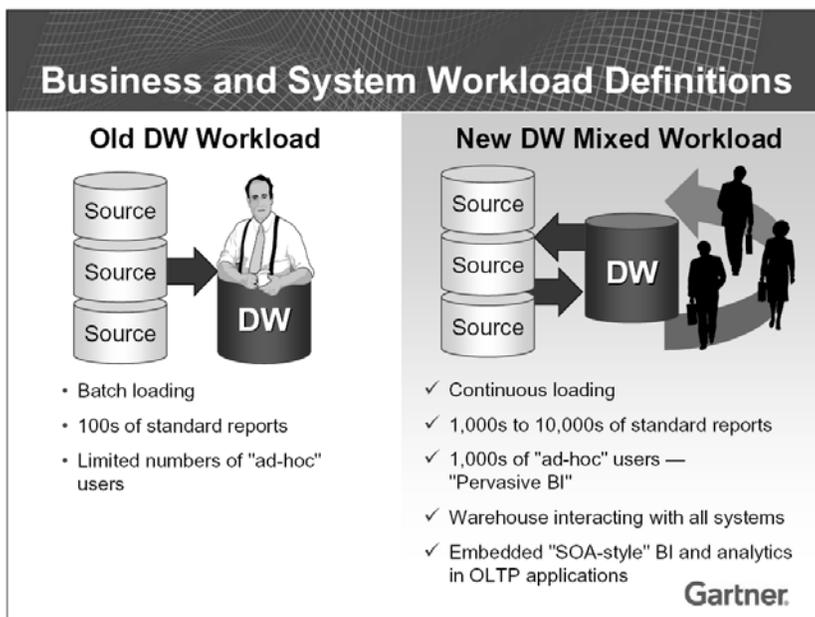
The real-time CDC solutions offer a continuous streaming "push" approach to delivering data. With these, data changes are captured as they occur and then immediately pushed to the target data warehouse or ETL system. Total latency can be brought down to minutes or even seconds, making that data near-instantly available to drive operational business intelligence and reporting.

### **Data Transformations – Where Do They Belong?**

As data warehouses evolve and become more operational with the benefit of real-time data feeds, the requirements for transforming the data has changed too. As described earlier in this paper, in the first three stages of data warehousing, data acquisition tends to be batch oriented. Data moves between relational and multidimensional structures and most of the transformations are handled at the data source, through scripts or on the chosen ETL engine.

As the data warehouse approaches real-time, transformations tend to take place in the data warehouse. This is often called an "E-L-T" approach: Extract, Load, then Transform. The data warehouse becomes more suitable to stage and transform the data in order to reduce data and analysis latency. This eliminates the need to aggregate changed-data on a centralized server until it is batch-processed and removes an intermediate step from the overall dataflow.

**Enterprise data warehouses are increasingly being designed to support these mixed workloads** so that the benefits of real-time data feeds can be fully realized and capitalized upon. Two leading vendor examples here are Teradata and the recently announced HP Neoview Enterprise Data Warehouse.

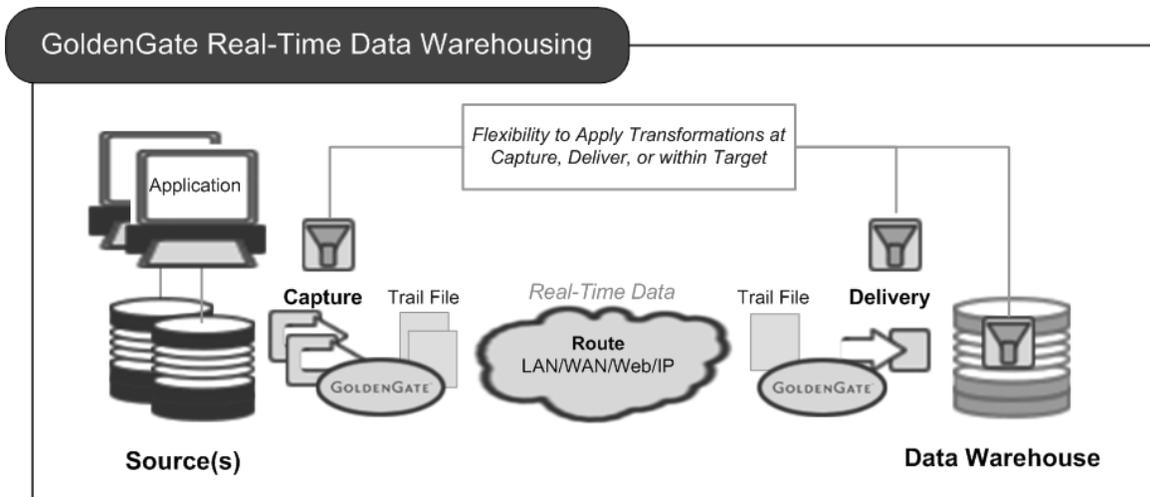


Source: Gartner 2007 Business Intelligence Conference, Mark Beyer

## 4. GOLDENGATE TDM FOR REAL-TIME DATA WAREHOUSING

To enable real-time data acquisition, GoldenGate's Real-Time Data Warehousing solution leverages its Transactional Data Management (TDM) technology to provide continuous, real-time capture and delivery of the most recent changed data between OLTP systems and the data warehouse. GoldenGate TDM implements transactional, real-time CDC in the "push" approach. As soon as new transactional activity is committed at the source system(s), that data is immediately moved to the data warehouse where it can drive enhanced, operational BI capabilities. Where certain heavy transformation requirements exist, GoldenGate may also send the data to a staging area in the data warehouse (to support an "E-L-T" architecture) or to the ETL server (see page 9 for more details on this).

In addition, GoldenGate's solution eliminates the need for batch window outages, is extremely low impact, supports the movement of large data volumes, improves the ability to recover data in the event of a failure or outage, and maintains the integrity of data transactions.



### How GoldenGate's Real-Time Data Warehousing Works

GoldenGate's Transactional Data Management (TDM) software platform includes process modules for capturing, routing, and delivering transactional data in real time across heterogeneous environments. For Real-Time Data Warehouse implementations, GoldenGate TDM provides exceptional capabilities:

**Data Capture** – A GoldenGate Capture module resides non-intrusively with the source database and continually captures any new transactions. The new data is immediately moved into GoldenGate Trail Files which stages the data in a database- and platform-independent universal data format. Trails not only enable heterogeneity but remove the risk of data loss or corruption, in the event of an outage at the source or target.

**Data Delivery** – GoldenGate continuously delivers all new data into the staging area of the data warehouse, with end-to-end latency in the sub-seconds. That means the most current data is always available for more advanced, agile business intelligence, actions, and reporting. And because smaller sets of data are being moved at any given time (unlike batch methods), there is very little overhead imposed on the source and IT infrastructure.

**Transformation Support** – GoldenGate provides native transformation support or the ability to call on business rules from third-party, external solutions. Transformations can be applied at the Capture or Delivery processes, or inside the warehouse when higher end-to-end performance is

desired. No additional middle-tier server is needed. GoldenGate's native transformation support includes table, row, and column selectivity, filtering, and mapping; a suite of date, math, string, and utility functions; and the use of user exits, stored procedures, and macros.

**Heterogeneity** – GoldenGate supports a wide range of database versions for Oracle, SQL Server, IBM DB2 OS/390 and UDB, Sybase, Enscribe, SQL/MP and SQL/MX, and Teradata running on UNIX, Windows, and HP Nonstop platforms, and can deliver data to any warehouse running on an ODBC compliant database. GoldenGate is the only solution for capturing data from Teradata, thus enabling dual-active solutions for Teradata customers.

**Flexibility** – Companies can quickly and easily involve new or different database sources and/or target systems to their data warehousing solutions by simply adding new GoldenGate Capture and Deliver processes. This allows GoldenGate to enable extended solutions for moving data back to the OLTP system for any closed-loop activities, or for simultaneously sending data to reporting instances, data stores, backups, or other target systems.

In addition, customers find that GoldenGate is significantly faster and easier to implement and maintain, compared to other data integration solutions such as ETL and EAI. GoldenGate Director™ provides a GUI-based add-on product for speeding the deployment, monitoring, and reporting on all the GoldenGate process modules supported across the enterprise.

### ***Combining GoldenGate TDM with ETL***

In addition to the data integration scenarios described above, GoldenGate can be implemented to support co-existence with ETL implementations – primarily using GoldenGate TDM to handle the real-time, continuous “data capture” or “E” part of the ETL process, but at a transactional level.

- **Continuous Feed to ETL Staging** – GoldenGate delivers the real-time data feeds to the target database's staging tables, where the ETL Server then picks up the data to apply transformations and finally loads it into the data warehouse's user tables.
- **Generate Flat Files for Micro-Batches** – GoldenGate's capture process can send the real-time changed data into a flat file format (via a customized user exit) that makes it easier to run micro batches throughout the day. Rather than the full ETL process needing to scan all the source tables, the control file generated by GoldenGate initiates the micro-batch process, and the data is then transformed and loaded as usual by the ETL product.

While drawbacks such as additional latency and the need to maintain the ETL server are introduced with these types of options, they may be considered if data transformation requirements are very complex. The benefit is that new transactional data changes are immediately captured with very low impact on the source OLTP systems, compared to the overhead imposed by ETL's extract process, and this makes mini, intra-day, or micro batches more achievable and efficient.

## **5. CUSTOMERS: REAL-TIME DATA WAREHOUSES IN THE REAL WORLD**

While the majority of enterprise data warehouses continue to move towards real-time data acquisition, there are numerous organizations that are seeing the benefits of real-time data warehousing today. These companies range from leaders in banking and financial services, airline travel, telecommunications, manufacturing, retailers, and ecommerce businesses.

These organizations continue to push the envelope with active data warehousing and are achieving measurable gains in their customer satisfaction levels and ultimately gains to their bottom line. Perhaps a significant investment into taking the enterprise data operational or active may not be justifiable for all businesses today – however as businesses becomes increasingly real-time, the enterprise analytics and data warehousing infrastructures should be built and ready to support faster business decisions.

Two excellent case study examples follow: Overstock.com, a large e-business retailer experiencing

exponential growth in data volumes; and Montefiore Medical Center, based in New York, which received an award from The Data Warehousing Institute for its use of GoldenGate to push real-time patient clinical data into warehousing and reporting systems.

### ***Overstock.com***

Overstock.com, Inc. is an online “closeout” retailer, offering discount, brand-name merchandise for sale over the Internet. The leading retailer offers customers the opportunity to shop smarter and more conveniently online for top-quality bargains. In 2004, Overstock.com’s revenue was up five-fold from 2002 and gross profits nearly quadrupled. With business growing at a rate of 100% a year, Overstock.com braced for continued growth and increasing customer demand. Experiencing 14 to 18 million hits a month to the website, Overstock.com recognized the need to both scale and streamline operations to better support its 24x7 customer transaction and reporting loads. Being an online retailer, there is zero tolerance for downtime. Overstock.com is always open and customers have come to expect that, so their IT infrastructure must be continuously available.

In addition, Overstock.com wanted to enable a real-time, single view of the customer to better understand purchasing habits, refine marketing efforts and more effectively drive business to the site. As a part of this initiative, Overstock.com decided to implement a Teradata Warehouse as well as other customer analytics applications from Teradata.

Overstock.com selected GoldenGate Software to move customer data from their Oracle databases supporting their retail site into the Teradata Warehouse running the Teradata database. Overstock.com has already started to see the benefits of leveraging customer data in real time. Their vice president of Data Warehousing Reporting and Analytics said, “If we send out an email campaign we need to know if it’s working. We need to know if consumers are clicking in the right place, if the email is driving consumers to the site, and if those consumers are making purchases. With GoldenGate we can access this kind of data in real time, rather than waiting one, two or three days. The speed is phenomenal and the integrity of each transaction is preserved.”

Since GoldenGate allows access to data in real time, Overstock.com can run reports around the clock, without putting additional strain on the operations system. In the past, the system could be tied up for long periods of time for a single data report, causing significant reduction in productivity.

“Accessing the data in real time using GoldenGate we can immediately see if we’re profitable and if our business processes are working. By using GoldenGate we now have the ability to track customer behavior as it happens and check our return on investment.”

In the past, Overstock.com was forced to treat all customers the same, whereas now the retailer can analyze customer behavior and purchase history to target marketing campaigns and service.

“We’ll be able to pinpoint which customers like the one dollar shipping campaign, or which customers prefer to shop in the discount section. We’ve never had the power to view our customers like this. With this information, we can provide better service to our customers. In the retail industry, it’s all about understanding your customers.”

With the rapid increase in purchasing demand, Overstock.com expects the new system to pay for itself within the first year, maybe sooner.

### ***Montefiore Medical Center***

The Data Warehousing Institute™ (TDWI) named GoldenGate customer Montefiore Medical Center, based in New York City, NY, a winner of the 2006 Best Practices in Data Warehousing Awards. Montefiore won the award in the “Right Time” Data Warehousing category for the organization’s innovative use of GoldenGate’s data integration solution to support real-time data warehousing.

Montefiore Medical Center, the University Hospital and Academic Medical Center for the Albert Einstein College of Medicine, ranks among the top one percent of all U.S. hospitals for investments in medical innovation and cutting-edge technology. Montefiore’s unique combination of “state-of-the-art” technology with “state-of-the-heart” medical and nursing care in a teaching and research environment offers its patients access to world-class medical experts, the newest and most innovative treatments and the best medical

center experience anywhere.

Montefiore was recognized for its pioneering approach to unlocking the value of historical, patient and medical data to improve decision support. GoldenGate's data integration solution is used to push critical patient data from their clinical information system (CIS) application, GE Healthcare's Centricity® Enterprise, and feed the data in real time to a Sybase data warehouse and other servers for reporting purposes. From there, report-writing servers running the Clinical Looking Glass™ (CLG), a custom-built decision support application, analyze patient data to help doctors and healthcare administrators make better, more informed decisions around the delivery of care to hospital patients. Montefiore also uses GoldenGate's solutions for uninterrupted application and service availability of its CIS and physician order entry systems, which is especially critical in a paperless environment.

"The stakes are high when you're dealing with patient data," said Eran Y. Bellin, MD, director of outcomes analysis and decision support for Emerging Health Information Technology (EHIT), a wholly-owned subsidiary of Montefiore. EHIT developed and continues to maintain the CLG decision support software system at Montefiore Medical Center, under the direction of Dr. Bellin.

"We believe real-time data warehousing for our Clinical Looking Glass system improves administration and decision-making, and ultimately helps provide a higher quality of patient care."

## 6. SUMMARY

Gone are the days when Data Warehouses were just for reporting, analytics, and forecasting. Today, more companies are working hard to make their data warehouses operational and active -- and thus more critical to the business. An important requirement is to enable lowest possible latency in which new data is delivered to the data warehouse, ideally in real time.

GoldenGate's solution for Real-Time Data Warehousing leverages Transactional Data Management technology that continuously delivers the most recent transactional data from OLTP systems with only sub-second latency.

Key benefits realized with GoldenGate:

- Real-time data for enabling more advanced, agile business intelligence
- Low-impact, high-performance solution for integrating data and applying business rules
- Does not require batch window outages
- Support for large data volumes and heterogeneity
- Can co-exist with existing ETL or "ELT" offerings, or even eliminate the need for a middle-tier ETL server
- Exceptional flexibility, easy implementation and maintenance
- Strong data recovery and data integrity capabilities

**About GoldenGate:** GoldenGate Software, Inc. is the leader in Transactional Data Management (TDM). Providing technology for capturing, transforming and moving transactional data in real time across major databases and environments, GoldenGate helps organizations implement a variety of critical solutions that deliver real-time access to real-time information across the enterprise. Major solutions are offered in the areas of Real-Time Data Integration and High Availability for transactional data systems.

With more than 400 customers worldwide, GoldenGate is used by industry leaders including Visa, Bank of America, Sabre Holdings, Comcast, DIRECTV, Federated Investors, Mayo Foundation, UBS, PayPal, AMD, and Overstock.com.

# GOLDENGATE<sup>®</sup>

## CONTACT INFORMATION

GoldenGate Software, Inc.  
301 Howard Street, San Francisco, CA 94105 USA  
Tel: +1 415-777-0200  
Email: [sales@goldengate.com](mailto:sales@goldengate.com)

For a list of worldwide offices and for more company information, please visit:

**[www.goldengate.com](http://www.goldengate.com)**

© 2009 GoldenGate Software, Inc.