

Delivering information you can trust

April 2008



IBM **Information Management** software

How IBM InfoSphere Change Data Capture can be an integral part of an event-driven architecture and service oriented architecture

Contents

- 2 Introduction**
- 3 Challenges**
- 4 Event-driven architecture: An overview**
- 5 IBM® InfoSphere™ Change Data Capture: An overview**
- 8 IBM InfoSphere Change Data Capture: Strategy**
- 8 IBM InfoSphere Change Data Capture: Architecture**
- 8 IBM InfoSphere Change Data Capture: Attributes**
- 11 Synopsis – Putting it all together**

The only constant is change, continuing change, inevitable change. Winning companies sense, anticipate and proactively respond to change, the rest trail behind.

Business Agility, the ability to anticipate and respond to business events in the form of change, opportunity or threat, is a key attribute in any organization's business strategy. And agility comes from aligning Business and IT "By Design", instead of drawing on coping mechanisms or "Going through the Motions".

Introduction

Event-driven companies – those that acquire, deploy, and profitably exploit real-time information, are most successful at sensing and responding to the changes or events that drive their businesses.

This calls for a redefining of existing traditional IT Architecture towards an Event-Driven Architecture. The term 'Event-Driven Architecture' refers to any application that reacts intelligently to changes in conditions, whether that change is a new customer registration, a termination of services from an existing customer, a hardware device malfunction; a power outage in one region of the country causing temporary shutdown or a sudden change in stock price. Depending on the size of the business, there are hundreds or thousands of notable events that occur every minute, every hour, and every day. By nature, some events may be positive, some negative, some may provide a business opportunity while others may pose a threat.

As a leader in real-time data integration, protection and auditing software, and as a pioneer in Event-Driven Architecture, IBM delivers one of the most robust, most heterogeneous and most high-ROI real-time solutions, all of which enable corporations to extract maximum value from mission-critical data and events.

Challenges

Events rule and events dictate – not only ongoing business actions but also future business strategies. And yet, most business processes have not been able to identify and react to them, until now.

With business phenomena such as competition, globalization and diversification and psychological phenomena such as consumer mentality, low-loyalty and low-tolerance, organizations need to get their act together to ensure customer retention, expansion and profitability.

Enter event-driven architecture, service-oriented architecture

Service Oriented Architecture (SOA) is becoming the defacto standard for technical infrastructure in organizations across all industries. It is an approach to building software applications as collections of autonomous services that interact without regard to each other's platform, data structures, or internal algorithms. SOA equips organizations in improving their ability to react to changing business dynamics and in leveraging their technical assets, thereby helping them gain competitive advantage.

2nd generation SOA, also referred to as Web 2.0, is the merger of SOA with Event-Driven Architecture (EDA). The outcome is the creation of new systems that exceed the sum of their parts.

EDA and SOA together allows companies to become what Gartner calls “Real-Time Enterprises” – enterprises that compete by using up-to-date information to progressively remove delays in the management and execution of its critical business processes. “Real-Time Enterprises” gain a competitive advantage with their data assets in motion, through new composite applications, real-time dashboards; KPIs and automated business process instantiation.

Event-driven architecture: An overview

Events are at the core of EDA. The way in which an event is recognized, enriched and flowed across a business, provides strategic competitive advantage. In order to become “Real-Time Enterprises”, organizations must be able to aggregate, integrate and react to events. Events are classified as:

Discrete data events

Discrete Data Events are those that are generated from a single system and represent granular activity in the business system. A shipping order, a change in the inventory level, or a new customer added to a customer relationship management system – all are Discrete Data Events; which are often analogous to a row in a table but may be the aggregate result of a few elements from multiple rows in multiple tables. What categorizes a data event as being discrete is its linkage to a single business entity such as a customer, or a product.

Composite events

Composite events are more complex in nature and represent business transactions aggregated from multiple sources such as order transactions from an inventory system combined with customer information like shipping address. Composite events are of high value to organizations because of their value in simplifying BAM, BPM, dashboards, workflow or composite application development.

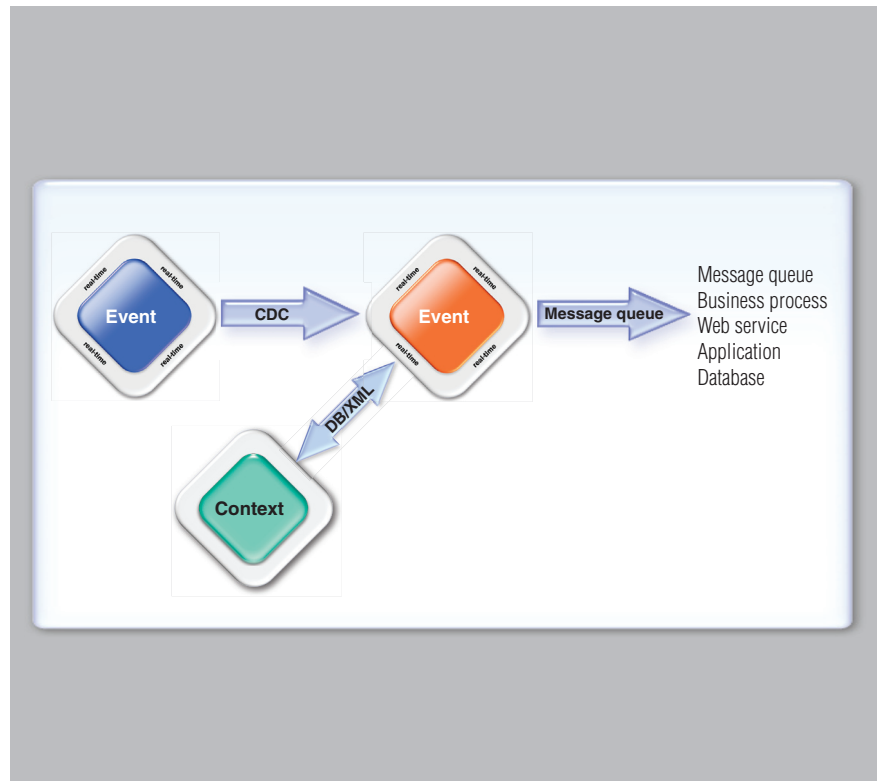
IBM InfoSphere Change Data Capture: An overview

As a business, the real challenge is getting business data which is at rest; to motion. Bridging the gap between where data is to where it will make a competitive difference, enables real-time insight into business operations to improve business agility and productivity.

This agility and productivity has to be achieved without falling into a Catch 22 situation where mobilizing data results in slower performance of key applications, higher bandwidth usage and higher implementation and operation cost.

Also, Real-Time decision support requires a higher level of data accuracy, data complexity, data availability and data freshness.

Figure 1: Overview of event-driven architecture



IBM InfoSphere Change Data Capture is a unique, end-to-end integration solution designed specifically to meet the technical requirements of transforming into a real-time enterprise. It leverages more than a decade of technological innovation in the area of Change Data Capture (CDC) technology; the groundbreaking technology which is a low-impact alternative for data integration. It eliminates redundant data transfer, thereby utilizing low network bandwidth. CDC technology, by reading log files generated by databases, captures data changes in the source application and transfers only the changed data to target applications. Reading database logs natively makes it feasible to process extremely high volumes of data without negatively impacting the performance of production systems.

IBM InfoSphere Change Data Capture is a pure and independent Java solution that can run on any J2EE compliant system, without requiring any special hardware or software. It is equipped with embedded EII functionality to create composite transactions that can target message queues, web services, applications, dashboards or databases, natively. EII is the integration of data from multiple systems into a unified, consistent and accurate representation geared toward the viewing and manipulation of the data. Data is aggregated, restructured and relabeled (if necessary) and presented to the user. Also, being GUI driven, IBM InfoSphere Change Data Capture is easy to implement, maintain, monitor, deploy and employ.

IBM InfoSphere Change Data Capture: Strategy

IBM InfoSphere Change Data Capture delivers event data in real-time to Message Oriented Middleware (MOM) – an effective solution for integrating business applications and legacy systems in heterogeneous environments. IBM InfoSphere Change Data Capture initiates a business process by capturing the relevant precipitating Business Event and delivers it in real-time into the MOM architecture. It delivers enriched events; events which are enhanced by additional data from other applications; in real-time, to the message oriented middleware technologies of leading Enterprise Application Integration (EAI), Business Process Management (BPM), and Service-Oriented Architecture (SOA) environments.

IBM InfoSphere Change Data Capture: Architecture

IBM InfoSphere Change Data Capture is founded on peer-to-peer architecture and creates a replication environment; which can be managed from a Graphical User Interface, a TSO command line, a DOS command line or through built-in APIs. The supported data sources are: UDB/DB2®, Oracle, Sybase, and Microsoft SQL Server. IBM InfoSphere Change Data Capture uses JMS (Java Message Service) to write to JMS Compliant Message Oriented Middleware. It provides the robust functionality required to manage and manipulate XML tree structure or hierarchy, enabling the user to develop custom structure that can be easily consumed by the application at the other end of the message bus.

IBM InfoSphere Change Data Capture: Attributes

The heterogeneous nature of business events demands technical requirements which by nature are flexible, risk-free and cost-effective. Unleashing the power of events and real-time data integration must neither introduce risk to existing infrastructure nor be rigid in its applicability. The solution must be flexible and adaptable across all business systems, regardless of the underlying database or operating system.

1. *Minimum risk – Inherently, IBM InfoSphere Change Data Capture is a reliable and secure solution that requires few changes to key applications or database structures. From the business paradigm, ensuring minimum risk requires identification and integration of only those events that directly impact agility, productivity and profitability.*
2. *Minimal impact – With its CDC technology, IBM InfoSphere Change Data Capture requires a minimum amount of CPU or memory, with relatively low impact on the performance of production systems or on organizational bandwidth. This is due to its log-based Change Data Capture (CDC) technology.*
3. *Seamless integration with existing infrastructure – IBM InfoSphere Change Data Capture supports event-awareness on existing infrastructure, without requiring new servers or databases. It is equipped and designed to work with multiple databases, multiple applications, across multiple platforms.*
4. *Information integrity – IBM InfoSphere Change Data Capture is equipped for the delivery and propagation of critical business events; and in the order in which they were generated by the source systems.*
5. *Ease of use – New technology means new learning and one of the major reasons for the failure, ineffectiveness or low ROI on new technology is due to its lack of intuitiveness to users. For any new technology to succeed and get results, it must be easy-to-use and easy-to-deploy, with features and functions accessible from an easy-to-use graphical interface. IBM InfoSphere Change Data Capture is designed for success through its easy use and easy deployment abilities.*

- 6. Real-time delivery to leverage information – IBM InfoSphere Change Data Capture is able to recognize and deliver events in real time, without major delays in delivering unless dictated by the business. There is neither a requirement for batch processing window nor any need to halt production systems to provide event-awareness capabilities.*
- 7. Source system agnostic – IBM InfoSphere Change Data Capture operates similarly and seamlessly across any database or platform used by the business system. It is able to recognize and aggregate complex or discrete events, regardless of the source being DB2® on Mainframe, Oracle on HP/UX or SQL Server on Windows.*
- 8. Target system agnostic – IBM InfoSphere Change Data Capture is able to serve events to whichever business system requires them, whether it is an application, database or messaging infrastructure.*
- 9. Multiple capabilities – IBM InfoSphere Change Data Capture may be utilized to address various different business issues and address various different business drivers. Its technology offers multiple utilities to provide a competitive edge to the business.*
- 10. Integrity, security – One of the greatest challenges for organizations today, is data integrity and security. Lack of trusted information can cost heavily to organizations and negate the advantages of real-time solution. IBM InfoSphere Change Data Capture ensures integrity of data and events and also ensures security of data and systems behind the events being propagated.*

Synopsis – Putting it all together

Business happens in real time and is always event-driven. Understanding, capturing and integrating events provides a competitive advantage through enhanced customer service, greater business agility and greater business visibility. The opportunity exists today to become event-driven and event-aware to gain a distinct competitive advantage. The technology exists today to capture and integrate business events, regardless of the complexity of the event or technical environment.

Knowing how best to utilize the capabilities of Event-Driven Architecture is going to depend upon the systems and processes in place today.

For more information

For more information about IBM Information Server, contact your IBM marketing representative or visit ibm.com/software/data/integration



© Copyright IBM Corporation 2008

IBM Software Group
Route 100
Somers, NY 10589

Printed in the United States
April 2008
All Rights Reserved.

IBM and the IBM logo, InfoSphere, DB2 are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries or both.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States, other countries or both.

Other company, product or service names may be trademarks or service marks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

***TAKE BACK CONTROL WITH* Information Management**