ORACLE®

Real-Time Data Warehouse using Oracle Data Integration Platform

Marco Ragogna EMEA Pre Sales – Data Integration

Agenda

- Data Integration Platform
- Real time Data Warehouse
- Oracle GoldenGate
- Success stories

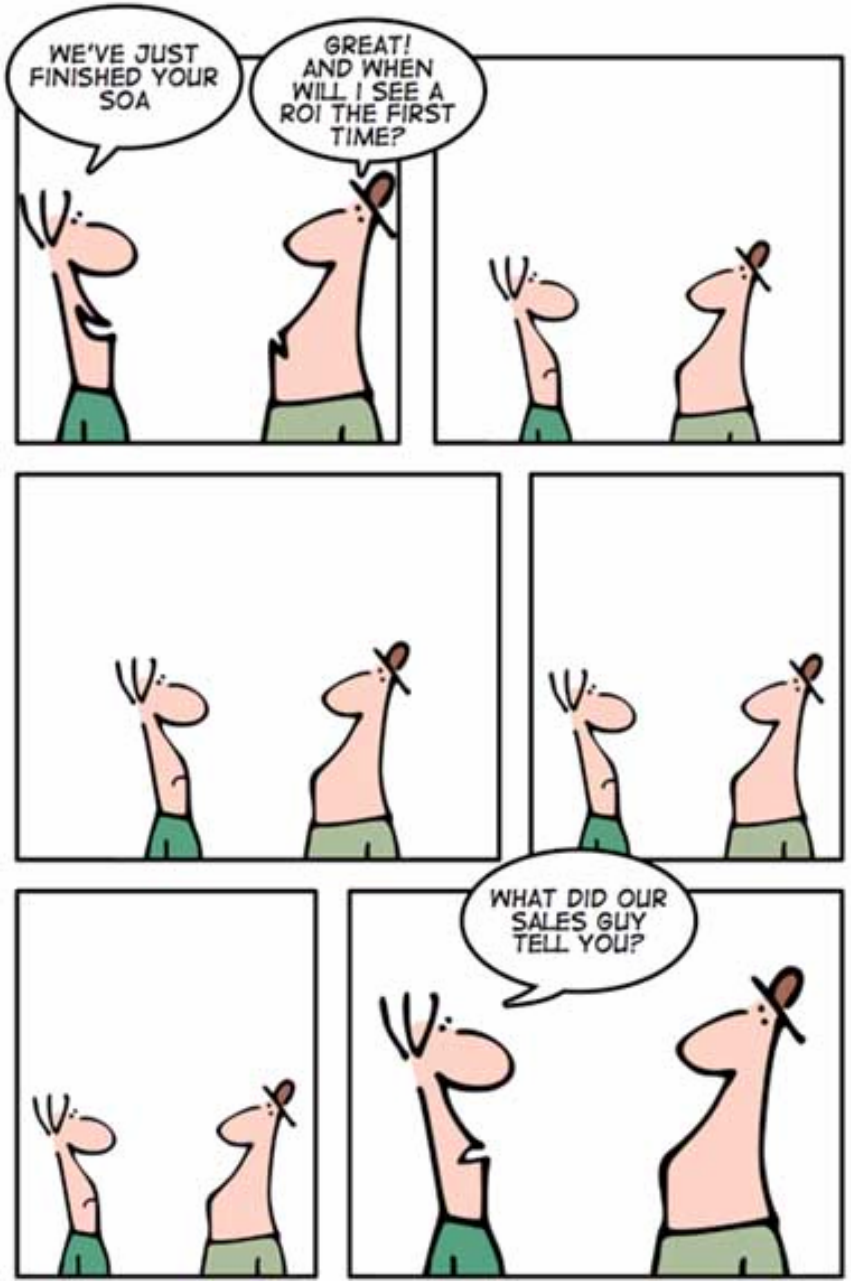

IT Obstacles to Unifying Information

What is it costing you to unify your data?

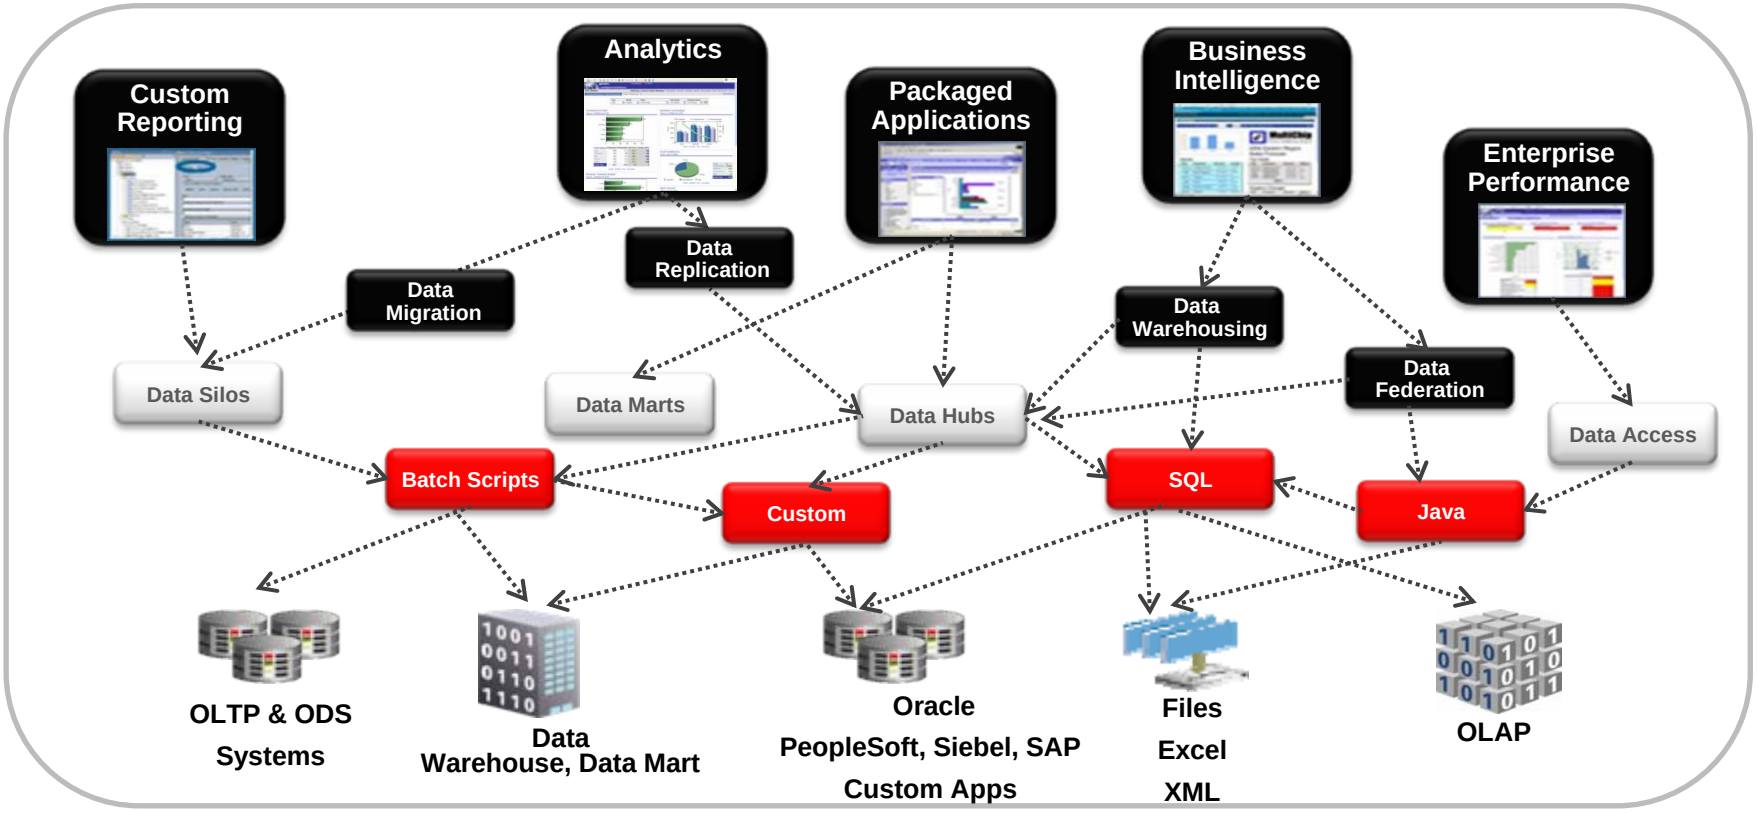

Fragmented Data Silos | **Slow Performance** | **Out of sync** | **Poor Data Quality** | **What's the cost?**

Oracle Data Integration

The solution for enterprise-wide real-time data

Databases

Distributed systems

Legacy systems

OLAP systems

OLTP systems

Mission critical systems and data

Business Intelligence, Performance Management

Data Warehouses, MDM

SOA

Dramatically improve the accessibility, reliability, and quality of critical data across enterprise systems

Oracle Data Integration Solution

Best-in-class Heterogeneous Platform for Data Integration

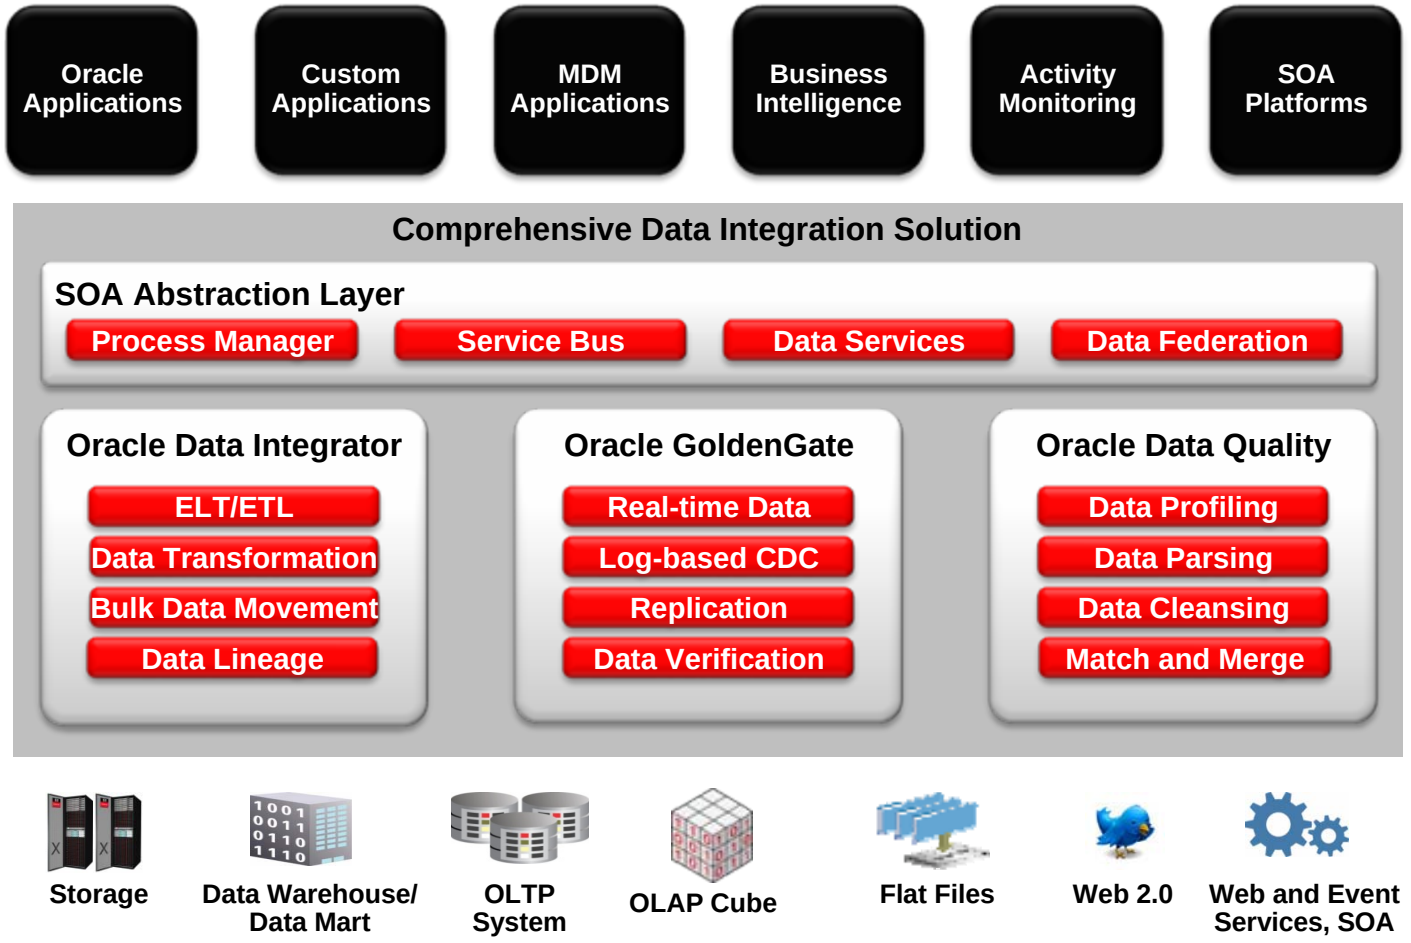

Audience for Oracle GoldenGate

How to make my DW support daily operational decision making?

1. DW Director
2. Existing DW needs to support business operations in addition to strategic decision making
3. DW Architect
4. Source systems need to be highly available– only 1.5 hours to load the DW
5. Starting a new DW implementation

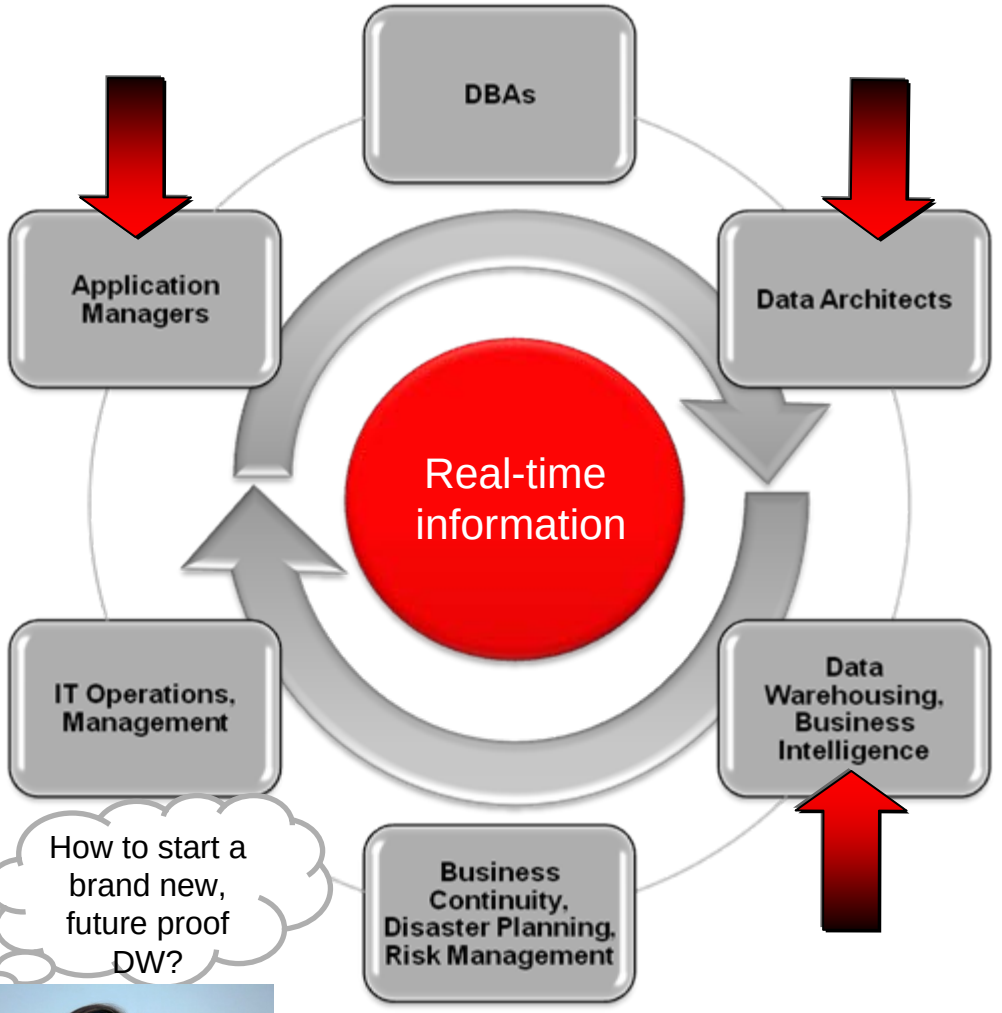

Will I be able to load all the data tonight?

How to start a brand new, future proof DW?

Architectures for Operational BI

Old approach

Run on the source OLTP system degrades performance of transaction processing

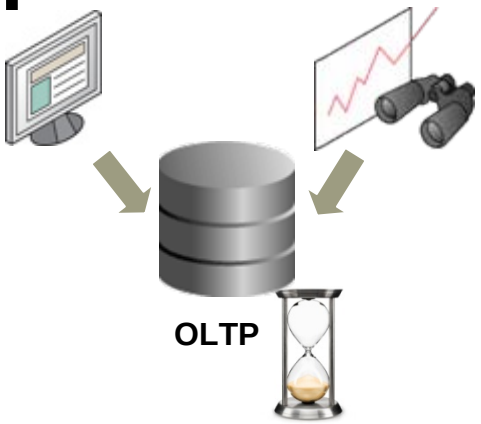

Best Practice

- Decoupling OLTP from Back end, then Back end from Front end
  - Several topologies can be useful

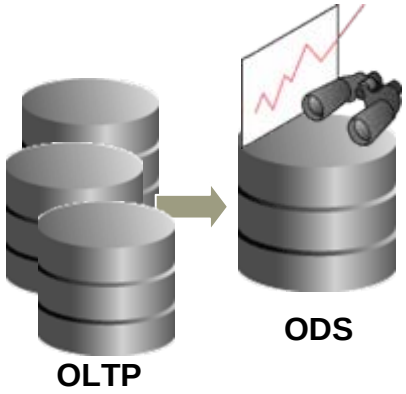

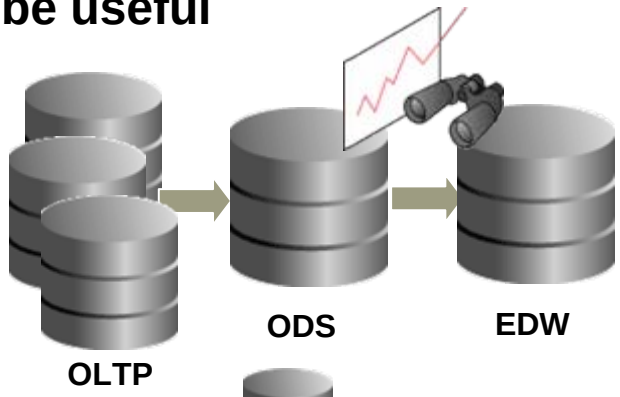

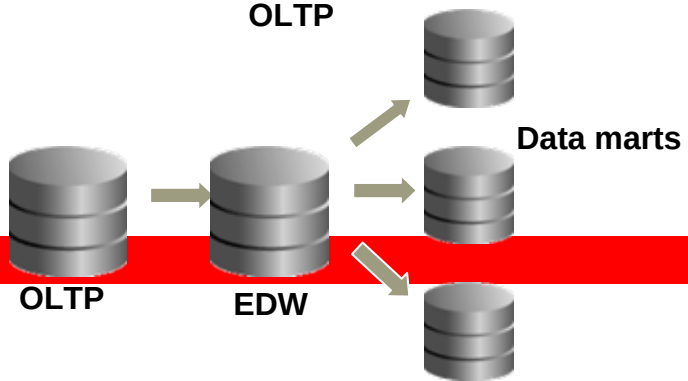

Common Methods for Data Acquisition

Overview

Common Methods

- Custom batch scripts

Shortcomings

- Not real time
- Requires batch windows
- Limited recoverability

- SQL queries
- Database triggers

- Near real time
- High overhead
- Intrusive
- Costly to deploy and maintain
- Very limited recoverability

Why Real Time?

Increasing the Value of Enterprise Information

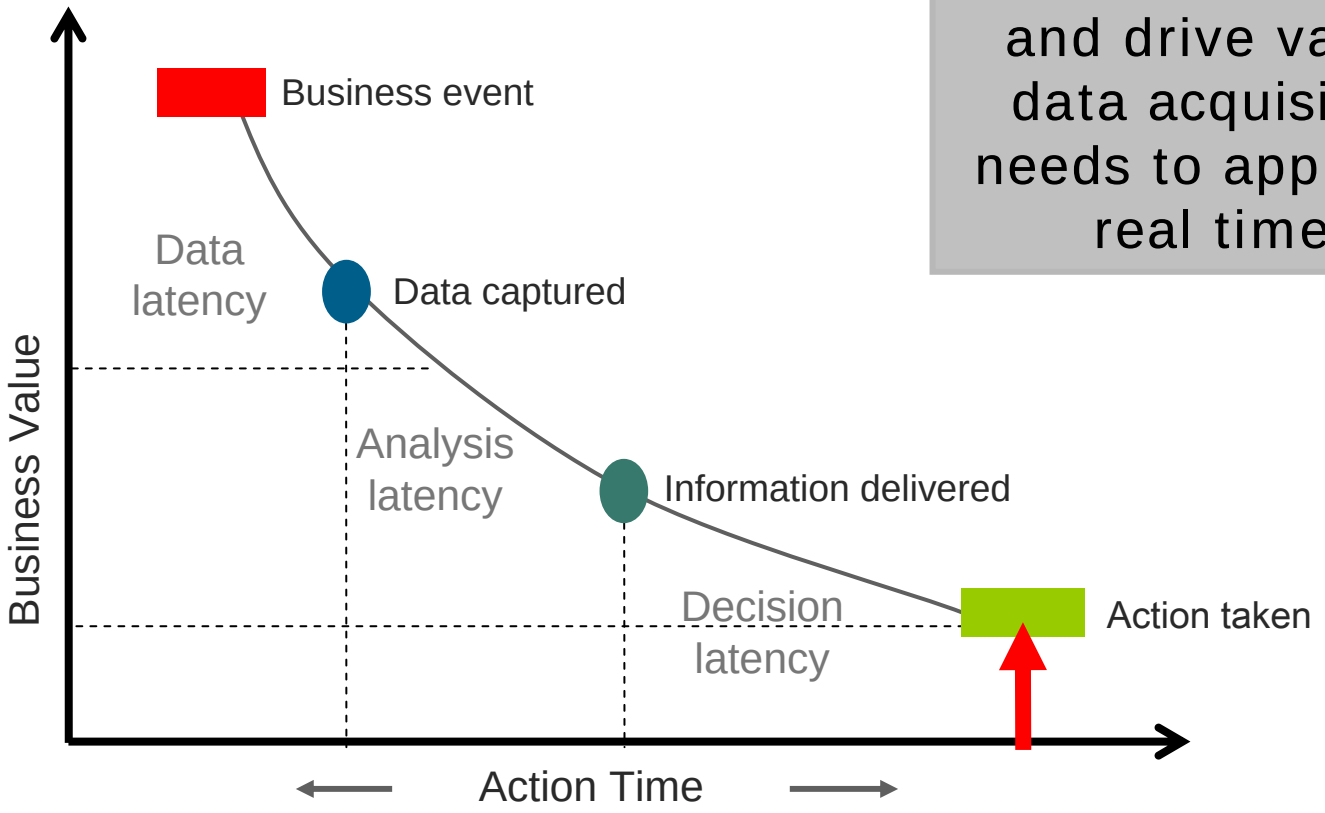

To reduce latency and drive value, data acquisition needs to approach real time.

From TDWI The Business Case for Real-Time BI
Based on concept developed by Richard Hackathorn, Bolder Technology

Example 1: Promotion Optimization-- Continued

- Results available next day

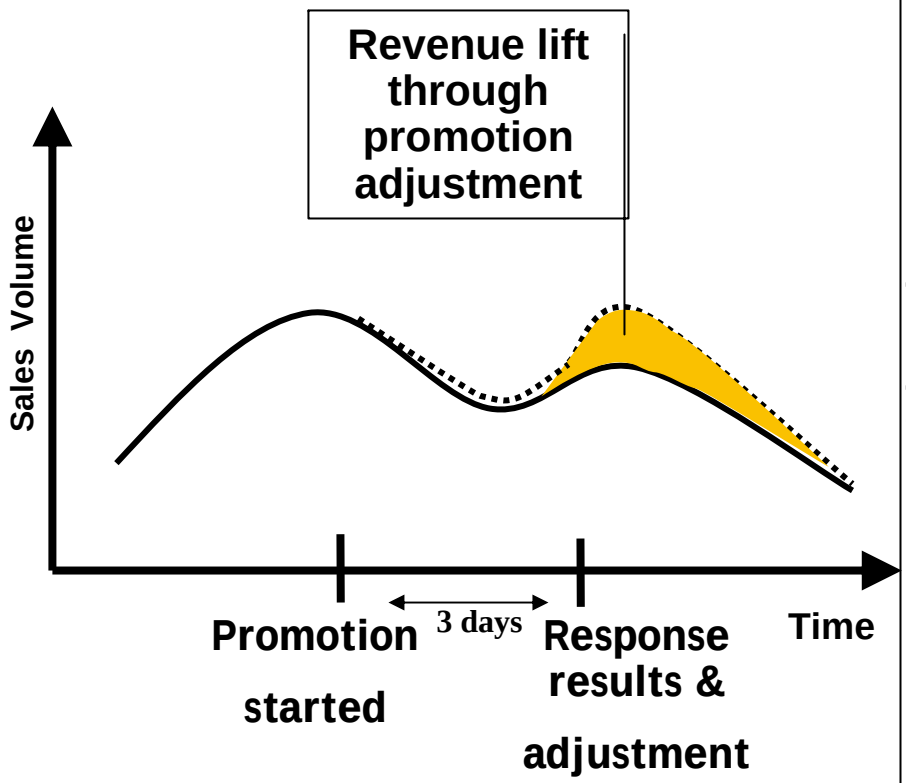

- **Results available in real time**

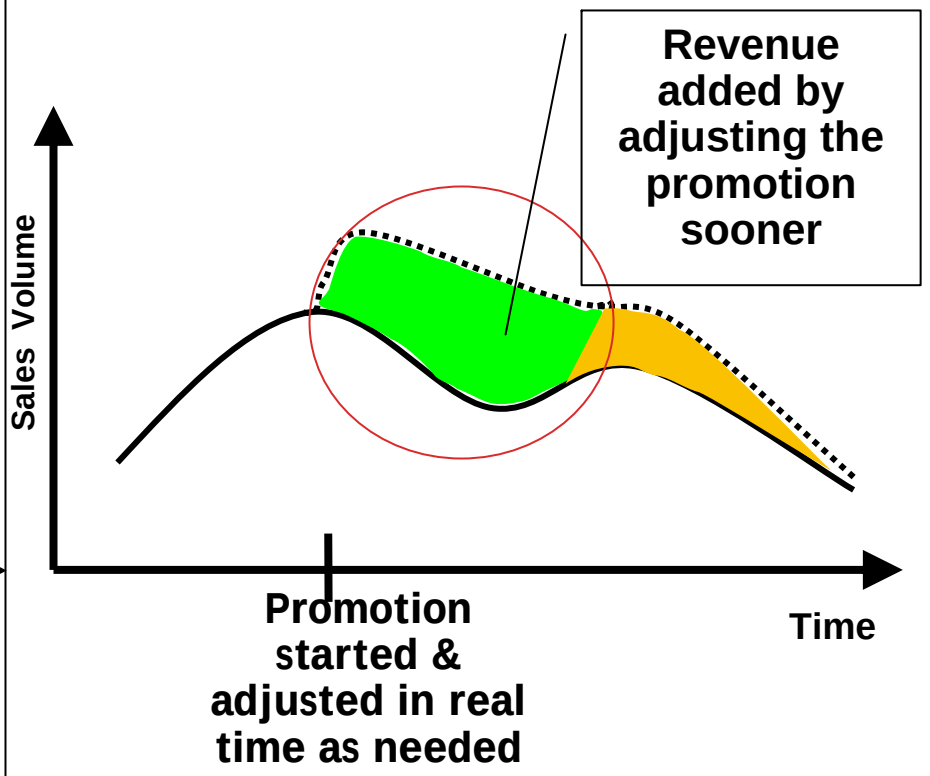

Historical sales ..... *Actual sales w/ promotion*

Based on concept developed by Richard Hackathorn, Bolder Technology

Today's Data Warehouse Needs an Alternative Approach to Meet the Changing Business Landscape

- Non-intrusive, real-time data feeds
- Timely data for analytical and operational systems
- Cleaner more authoritative data
- Global data distribution and synchronization
- Upgrade/migrations without interrupting operations

Data Integration for DWH: change to ELT, real time

- Data Warehouse Loading

For Typical Data Warehouse, Use Data Integrator Enterprise Edition

- ODI-EE has GUI design environment for ETL developers, basic changed data capture for typical ETL situations, rich transformations and ability to add-on Data Quality

Add GoldenGate to ODI-EE

- When sensitivity about source system invasiveness is high, very low latency is critical, transactional integrity in the data warehouse is important

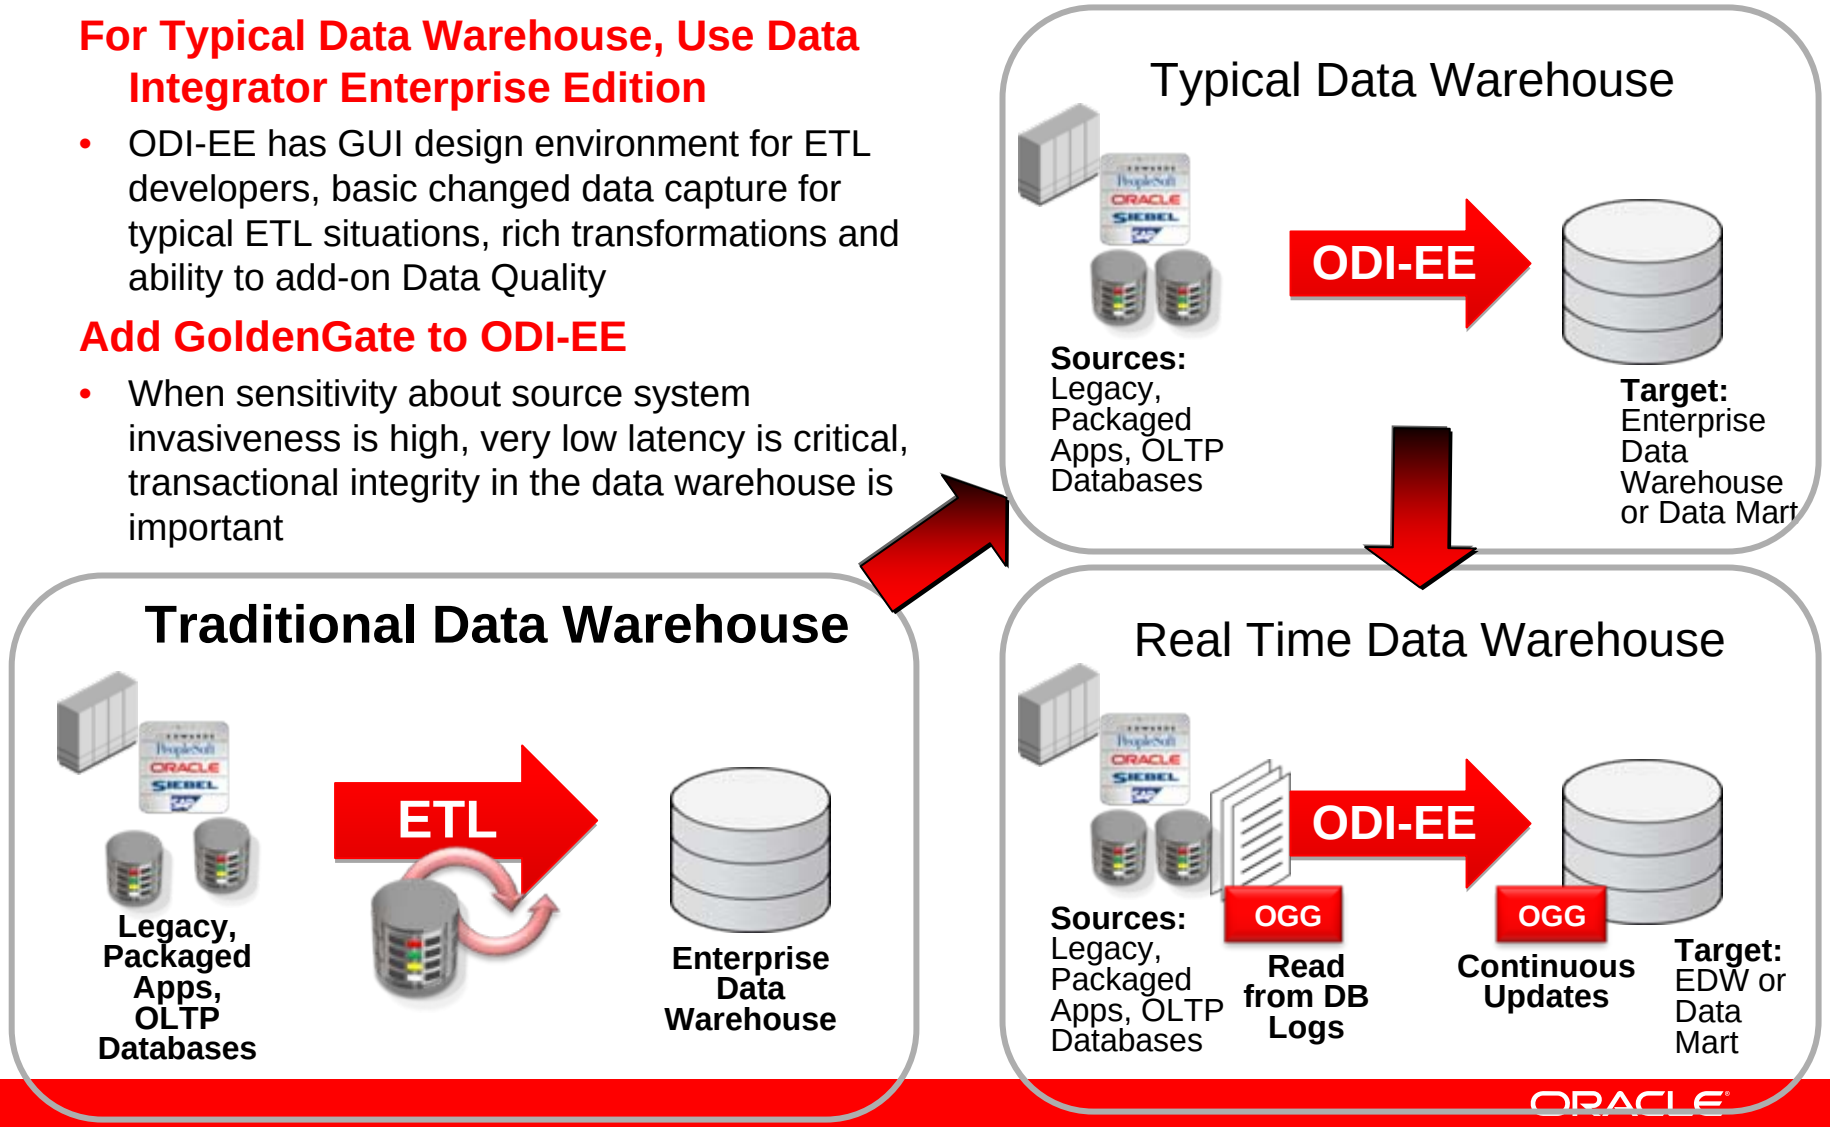

Oracle GoldenGate

Real-Time, Heterogeneous Change Data Capture and Replication

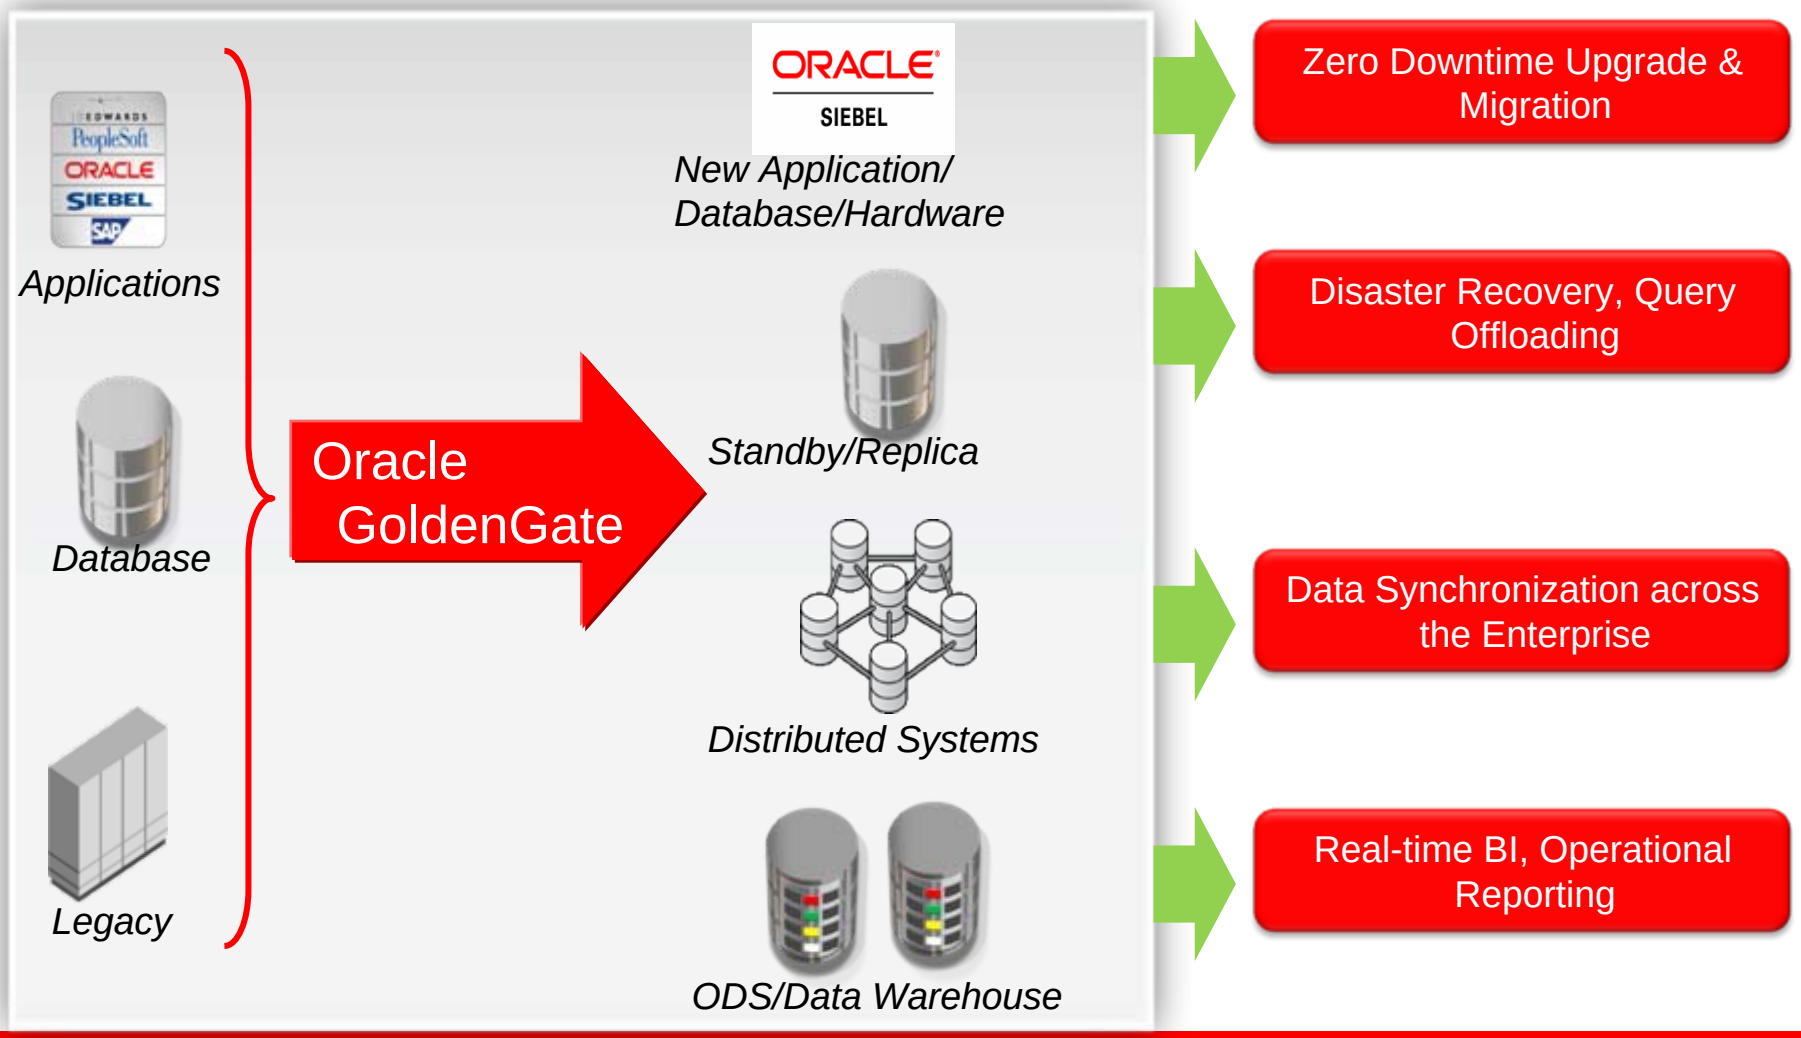

What is Oracle GoldenGate?

Oracle GoldenGate provides low-impact capture, routing, and delivery of transactional data across heterogeneous environments in real time

Key Differentiators:

Performance	Non-intrusive, low-impact, sub-second latency
Flexible and Extensible	Open, modular architecture - Supports heterogeneous sources and targets
Reliable	Maintains transactional integrity - Resilient against interruptions and failures

Oracle for Real-Time Data Warehousing

Fastest Real-Time Data Integration, Fastest Bulk Data Transformation

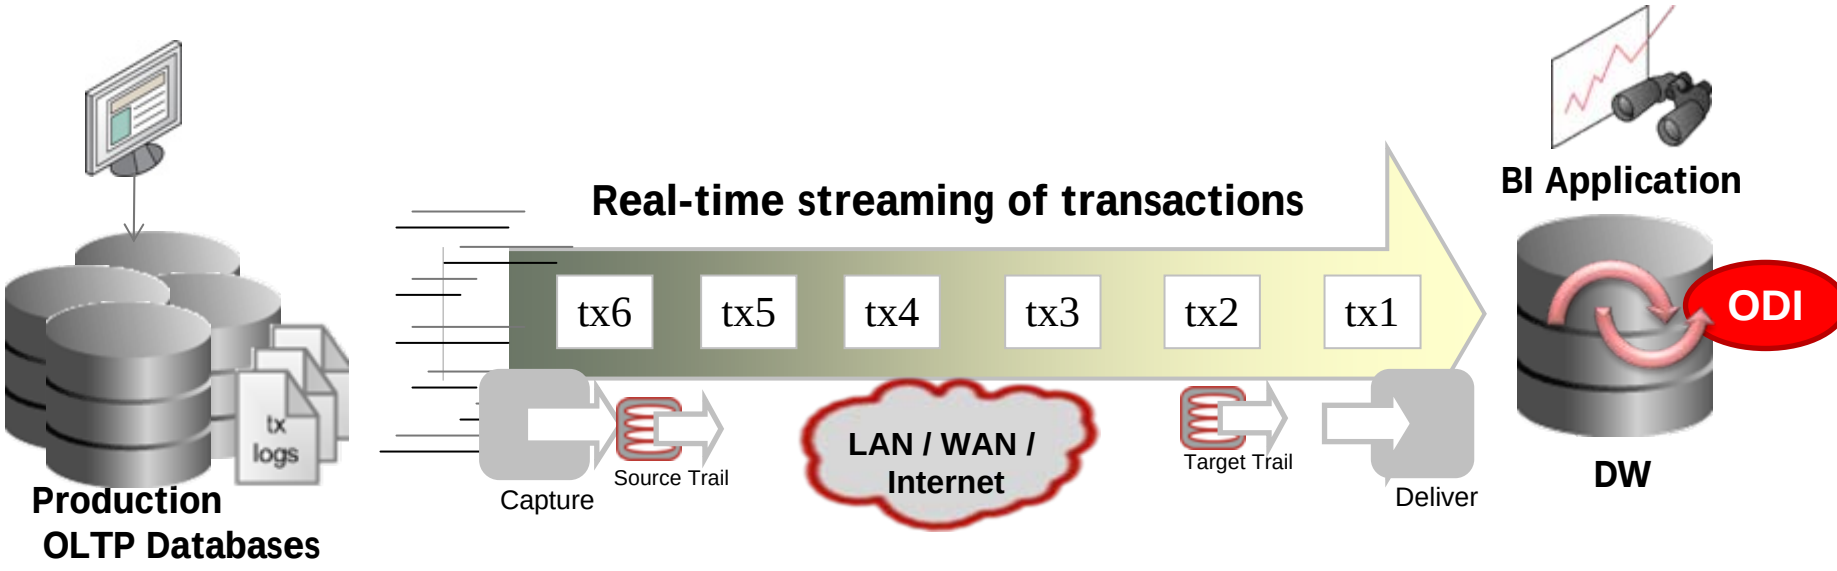

- Sub-second data latency
- Minimal overhead and no batch windows
- High-performance, in-database transformations
- Read-consistent changed data with referential integrity
- Complete data recoverability via Trail files

How Oracle GoldenGate Works

Capture: committed transactions are captured (and can be filtered) as they occur by reading the transaction logs.

Trail: stages and queues data for routing.

Pump: distributes data for routing to target(s).

Route: data is compressed, encrypted for routing to target(s).

Delivery: applies data with transaction integrity, transforming the data as required.

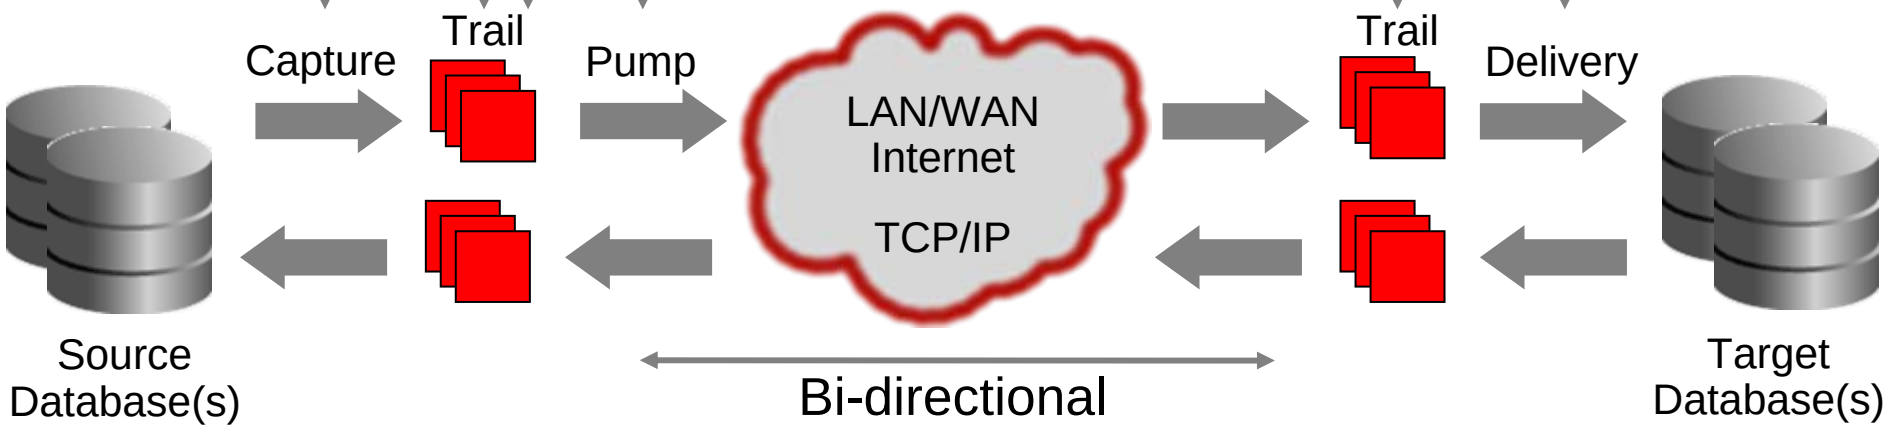

GoldenGate: Key Architecture Highlights

- Distributed Loosely Coupled Process Based Architecture
- Non-Quiesced Instantiation
- Integrated use with ODI
- Extensibility: Change Data Capture API (VAM)
- Customization: User Exits, SQL callouts, Macros
- Asynchronous Change Propagation
  - Committed Transactions Only
  - DML, Metadata, DDL
- Canonical Trail Format
  - Interoperable across Heterogeneous Systems
- Data Transformation Rules and Filtering (table, row, column)
- Batched Operations, Transactions

GoldenGate Capture

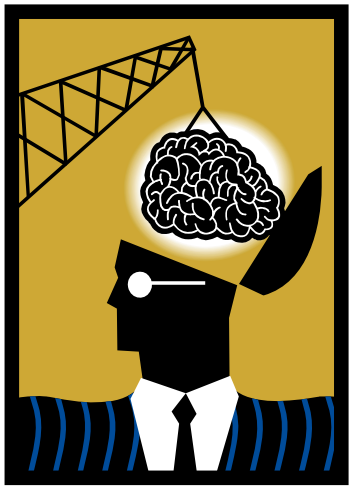

- Log-based change data capture
  - High volume
  - Low overhead
- Decoupled architecture
  - Multiple capture processes may be used to scale, but generally not required
  - Possible to split “hot” tables into a separate capture process
- Filtering and Compression
  - Unnecessary change data if discarded immediately
  - Updates and deletes are “compressed” by default
  - Before images are discarded by default

Oracle GoldenGate Capture

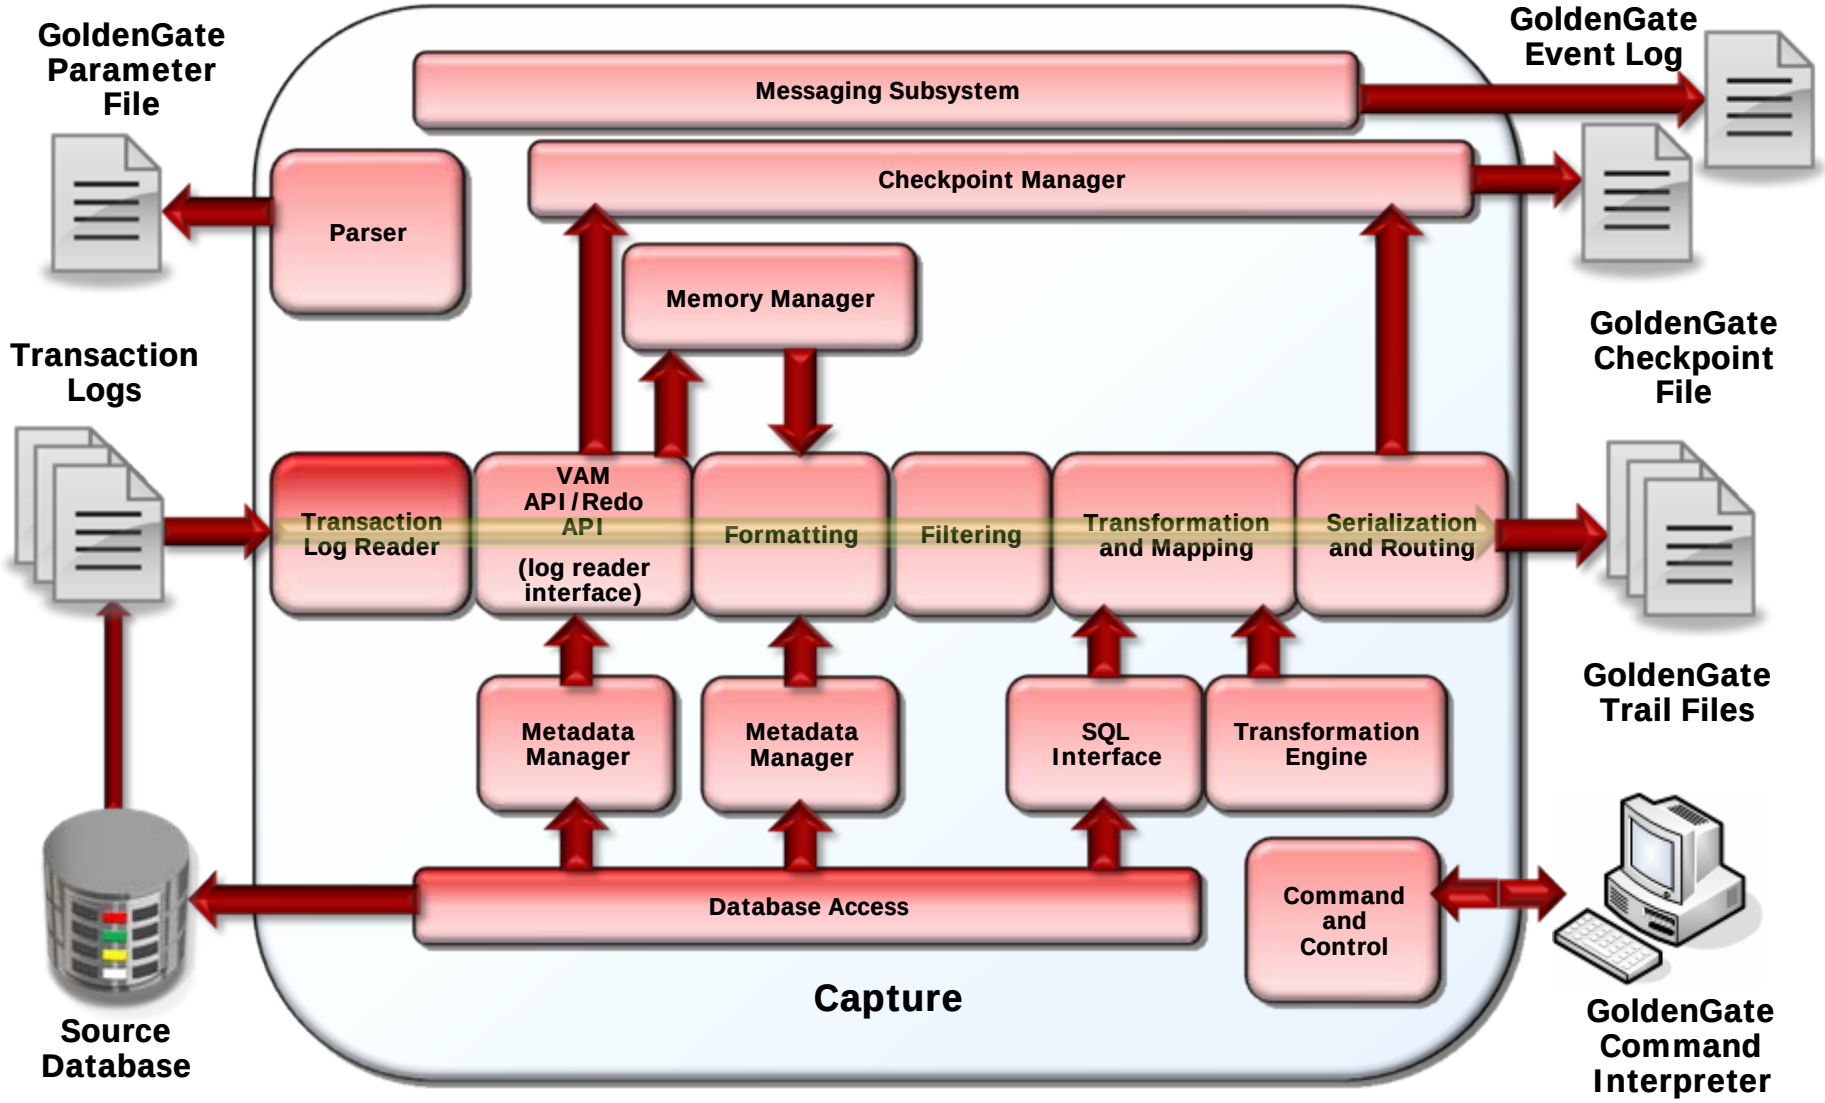

GoldenGate Checkpointing

- Capture, Pump, and Delivery save positions to a checkpoint file so they can recover in case of failure

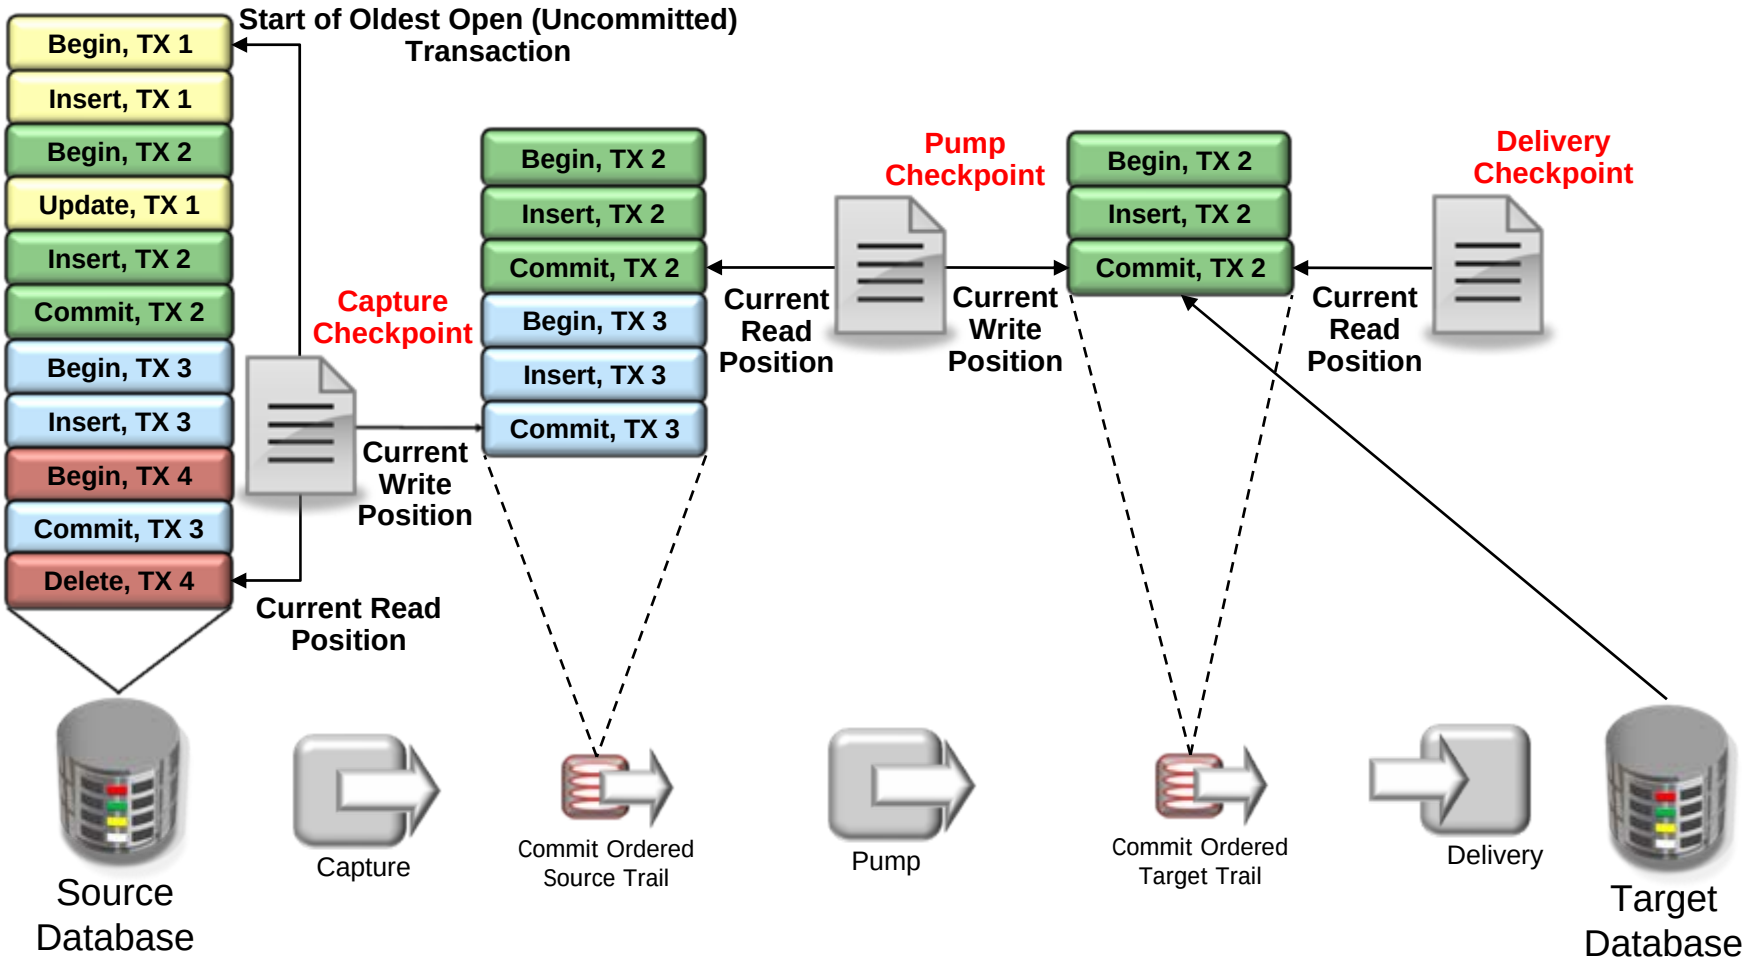

GoldenGate Delivery

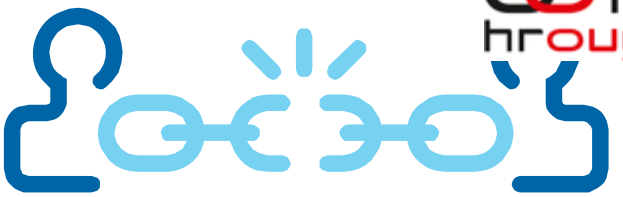

- Decoupled architecture
  - Multiple delivery processes may be used to scale
  - Possible to split “hot” tables into a separate delivery process
  - Possible for multiple delivery processes to split the work for a single table
- Transaction grouping
  - Small transactions are grouped by default to reduce commit overhead
- Record batching
  - Records are batched by table and operation within a single SQL execution
  - Automatic reordering of batches

GoldenGate Instantiation

- Start delivery at or after given backup CSN (245 in this case)
  - Current source Commit Sequence Number (CSN) is 356
- Once Delivery is current the source and target are in sync

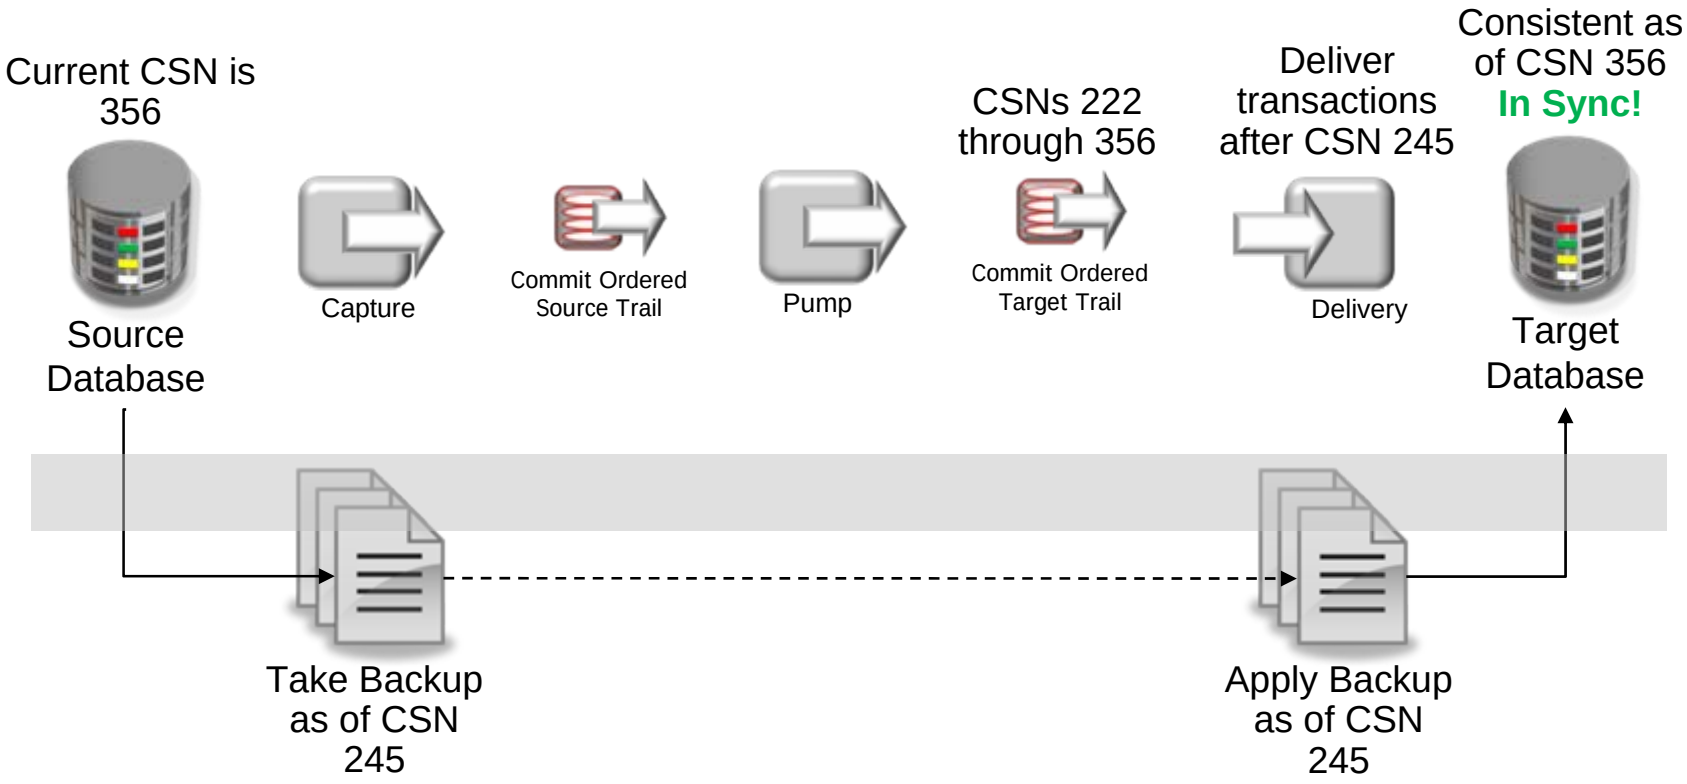

GoldenGate – Scaling for Performance

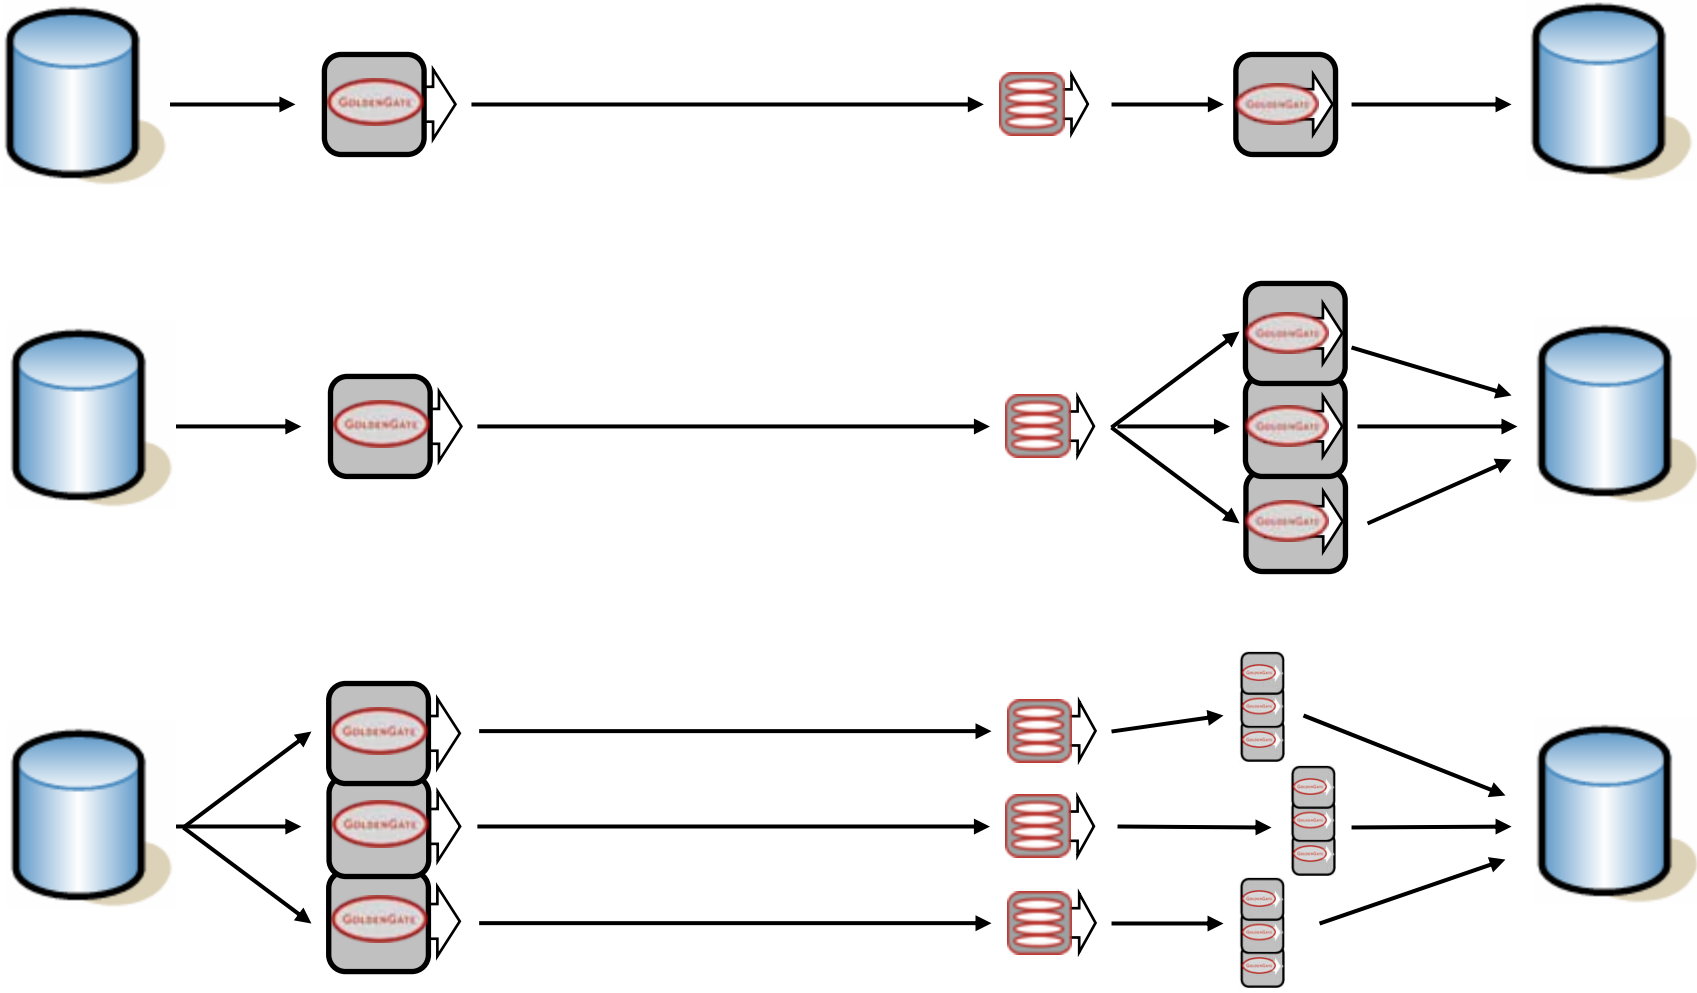

GoldenGate Heterogeneous Platforms

Databases	O/S and Platforms
Capture: ▪ Oracle ▪ DB2 ▪ Microsoft SQL Server ▪ Sybase ASE ▪ Teradata ▪ Enscribe ▪ SQL/MP ▪ SQL/MX Delivery: ▪ All listed above, plus: ▪ Oracle Exadata ▪ HP Neoview, Netezza, Greenplum, and any ODBC compatible databases ▪ ETL products ▪ JMS message queues ▪ MySQL ▪ TimesTen	**Windows 2000, 2003, XP** **Linux** **Sun Solaris** **HP NonStop** **HP-UX** **HP TRU64** **HP OpenVMS** **IBM AIX** **IBM z/OS**

Differentiator: Heterogeneity

Oracle Capture

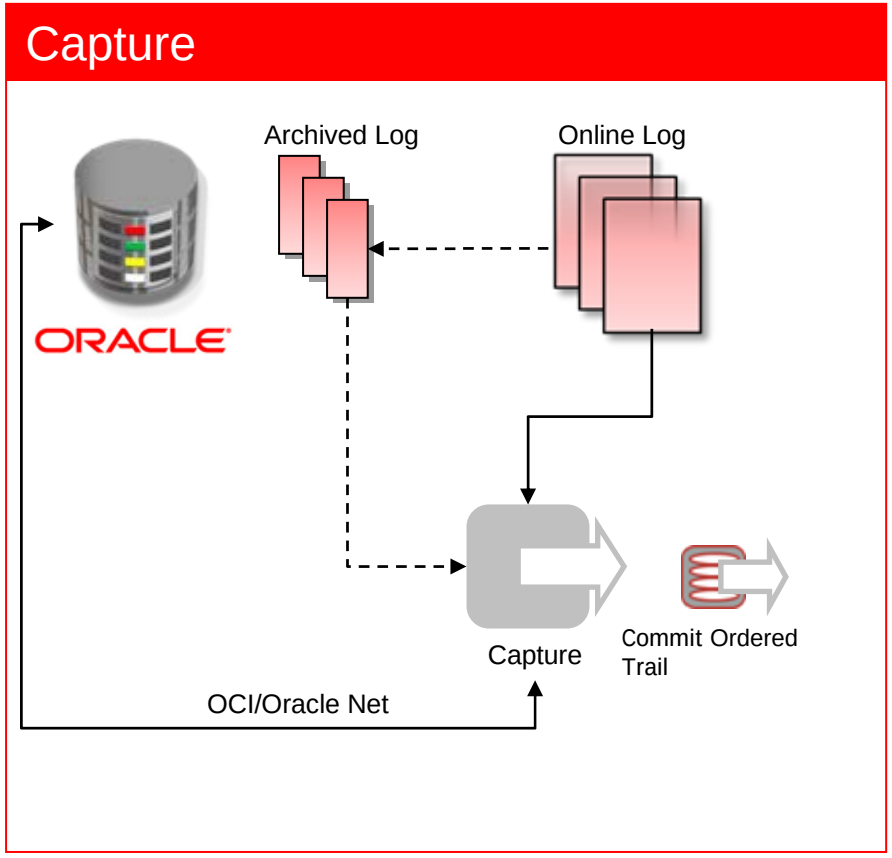

- Support for Oracle RAC
- Support for ASM
- Support for index-organized tables (IOTs) with overflow
- Support for clustered tables
- Support for object tables
- Support for object types (UDTs)
- Support for DDL operations
- Archived log only (ALO) mode
- Off-platform capture (LOGSOURCE)
- Multi-threaded capture

Other Key Features

Delivery to Flat File

- Delimited/fixed format data file and control file

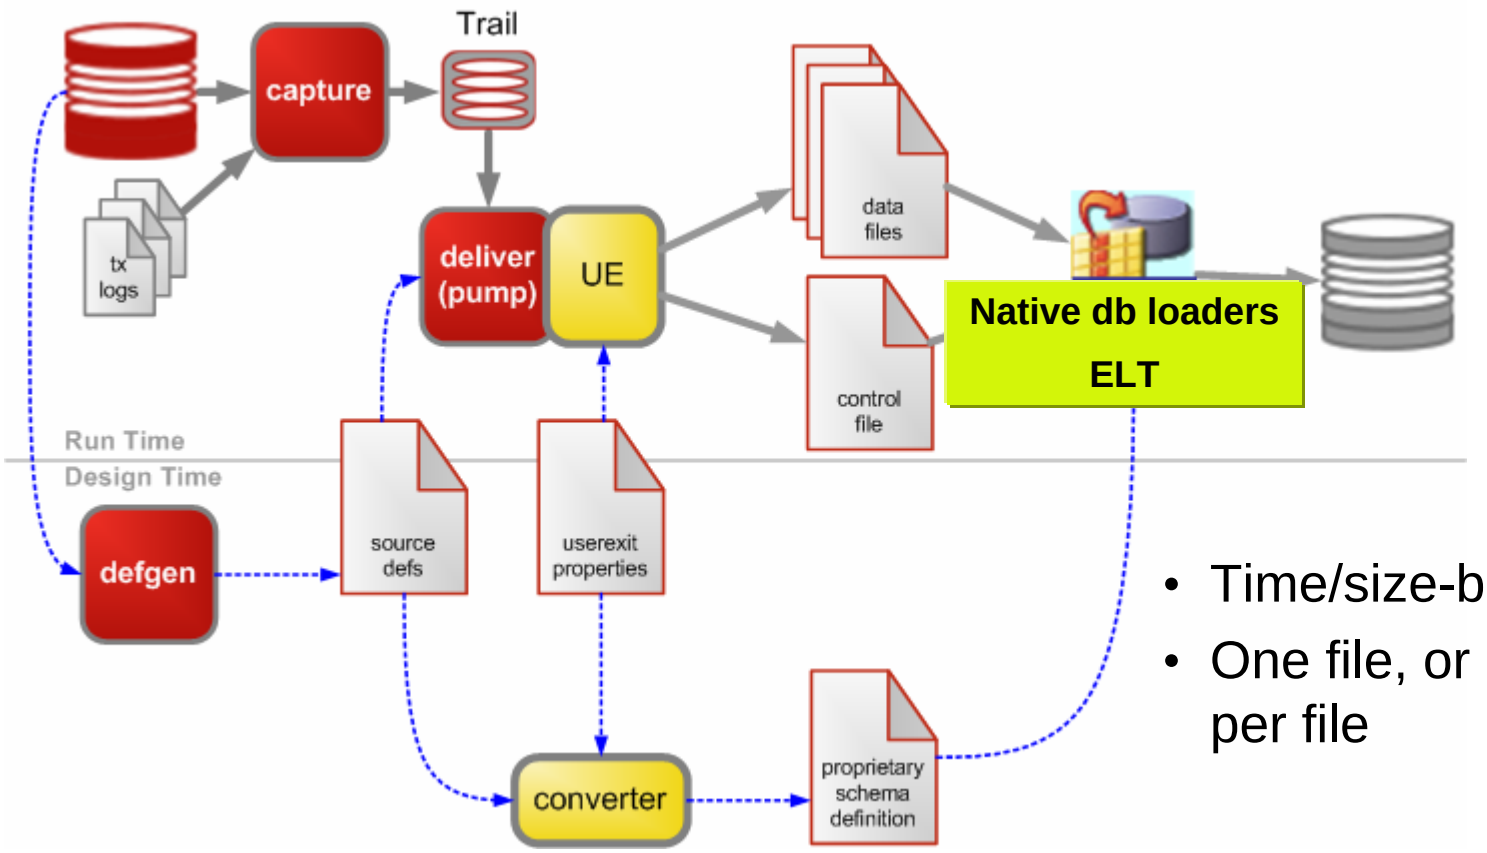

- Time/size-based rollover
- One file, or one table per file

Oracle GoldenGate Topologies

Unidirectional
Query Offloading

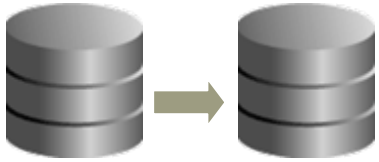

Bi-Directional
Live Standby or
Active-Active for HA

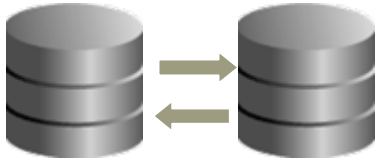

Peer-to-Peer
Load Balancing,
Multi-Master

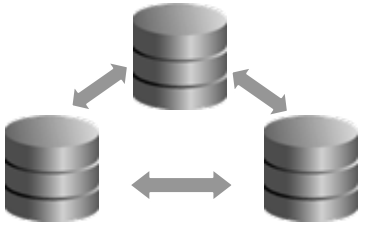

Broadcast
Data Distribution

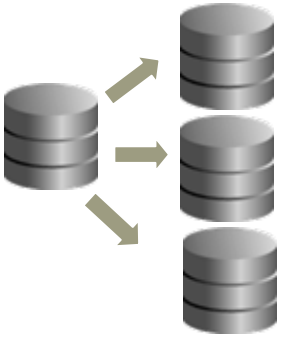

Integration/Consolidation
Data Warehouse

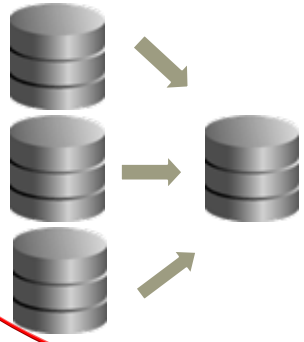

Cascading
Data Marts

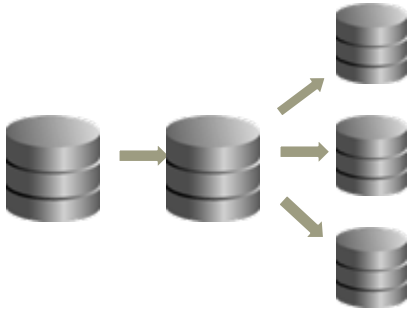

Zero Downtime Oracle Upgrade Implementation Steps:

Example of 9*i* → 11g Cross-Platform

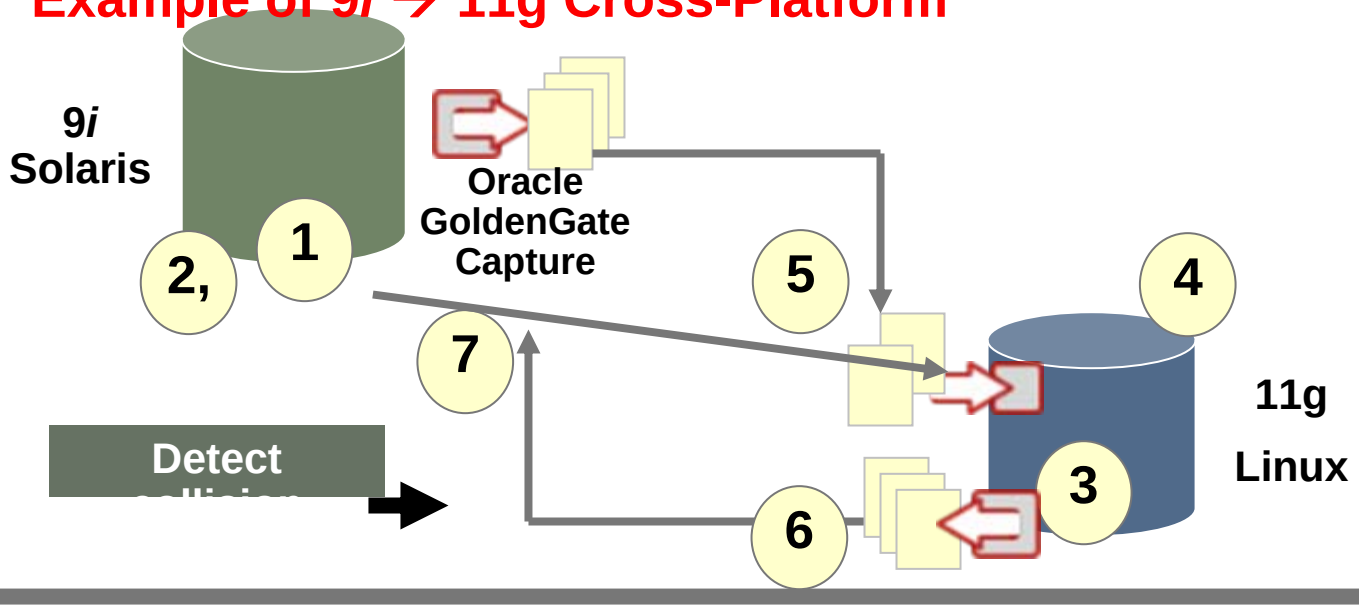

1. Start Oracle GoldenGate Capture module
2. - 4. Initial loading, export import of a new 11g target db (ELT/flat files/jdbc/native db loaders/import export tablespaces etc.)
5. Start Oracle GoldenGate Delivery module at target
6. Start Oracle GoldenGate's Capture at 11g
7. Start Oracle GoldenGate's Delivery process 9*i* (old source, contingency)

Oracle Real Time Business Intelligence

Complete, Flexible, Integrated

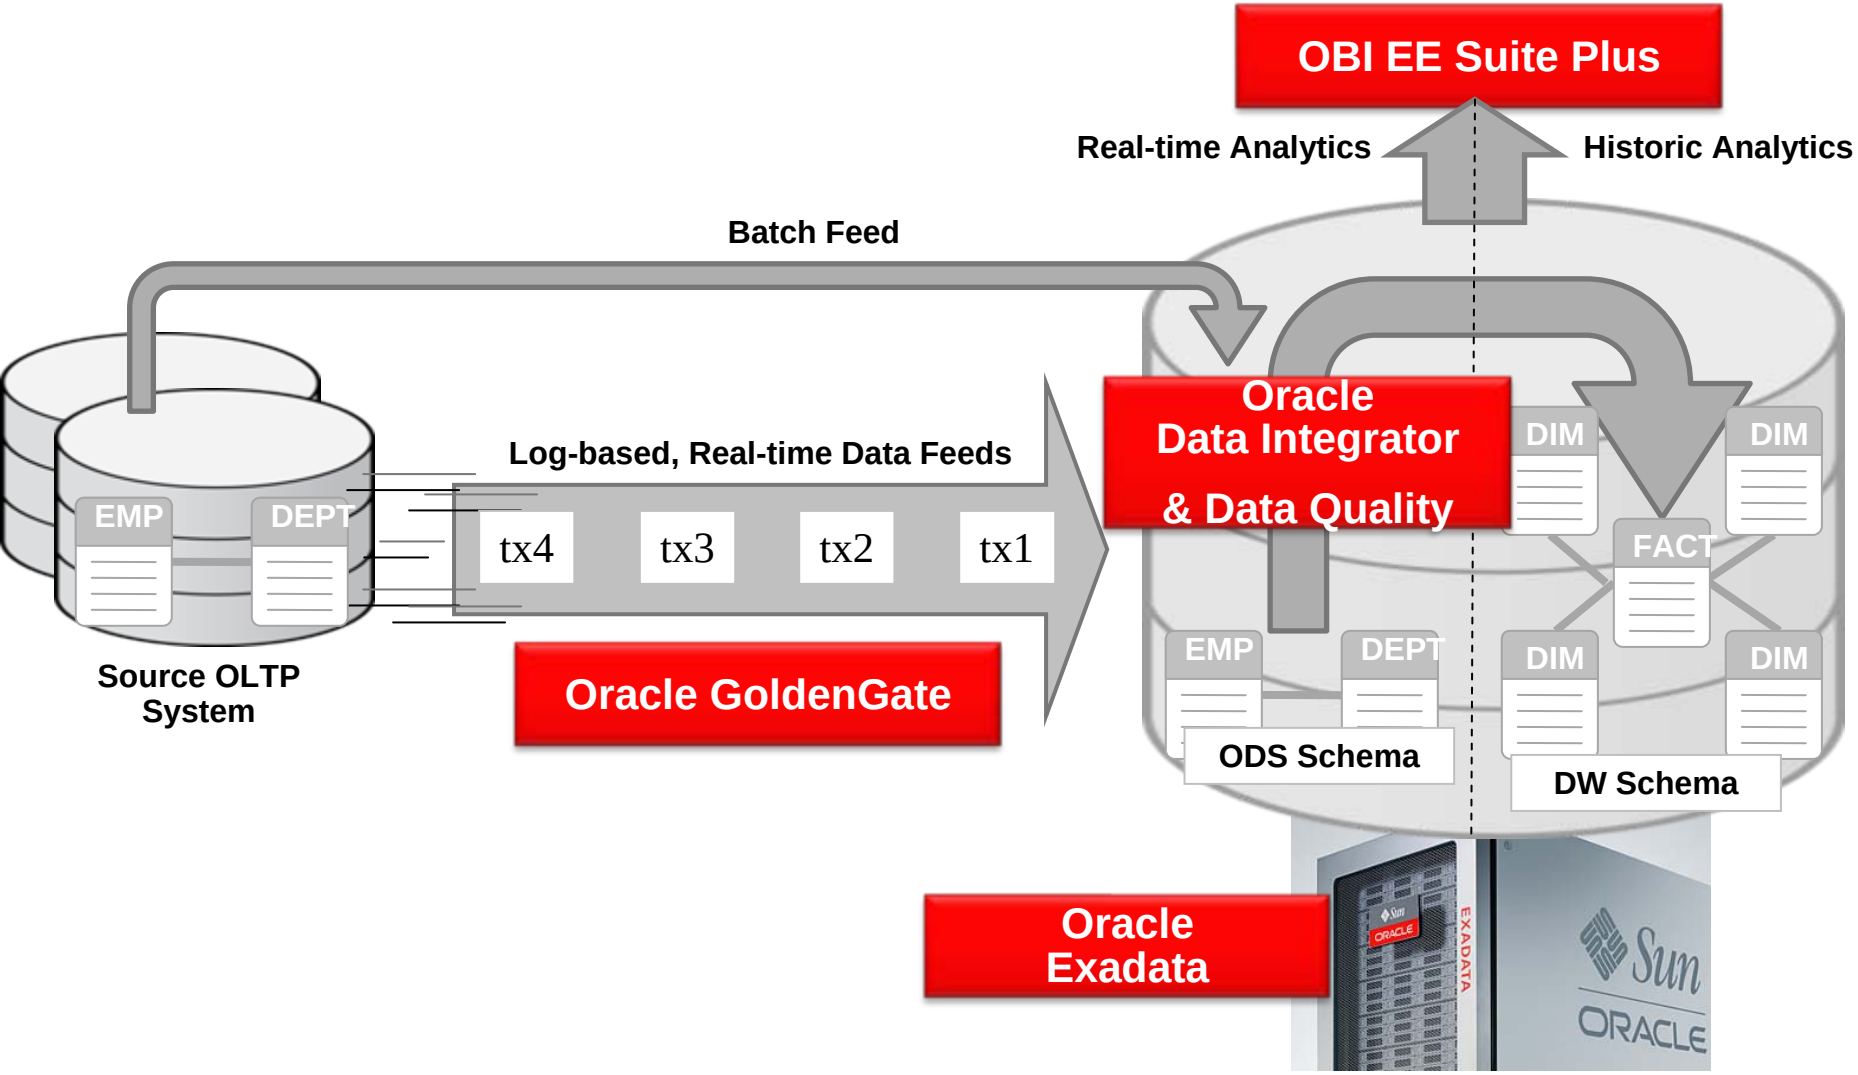

Differentiator: E-LT Architecture

High Performance

Transform in Separate ETL Server

- Proprietary Engine
- Poor Performance
- High Costs
- older approach

Transform in Existing RDBMS

- Leverage Resources
- Efficient
- High Performance

Benefits

✓ Optimal Performance & Scalability

✓ Easier to Manage & Lower Cost

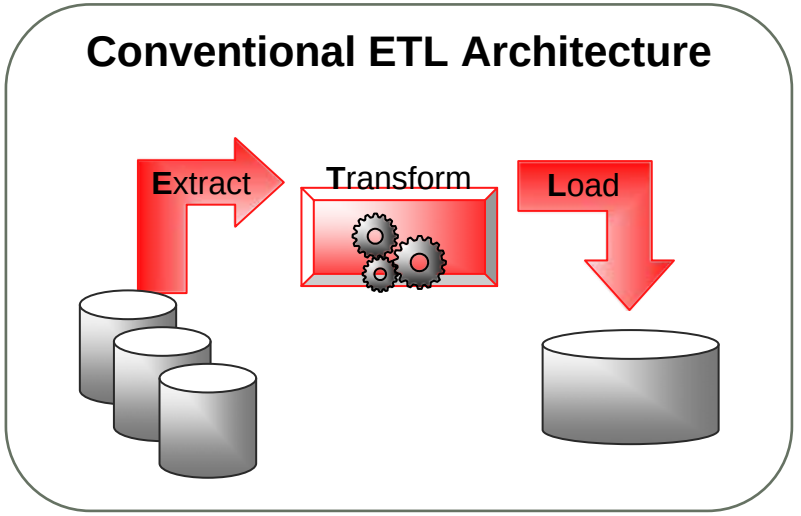

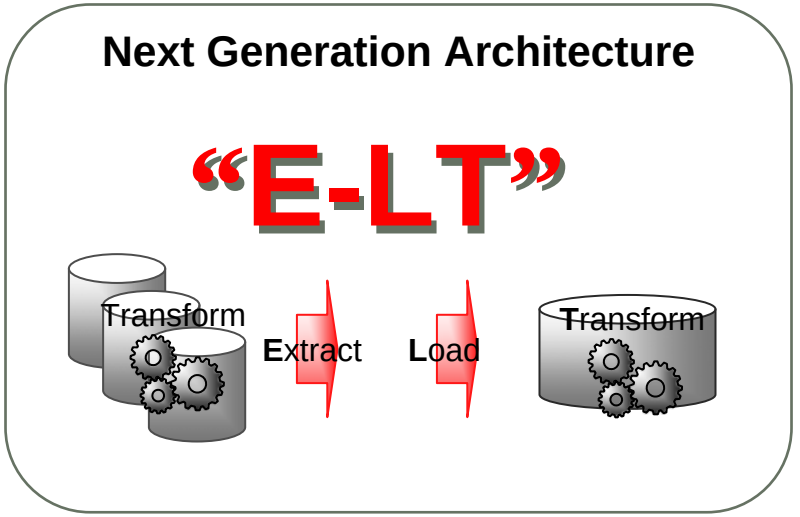

Traditional procedural ETL

Traditional ETL row to row complexity

Informatica Mapping

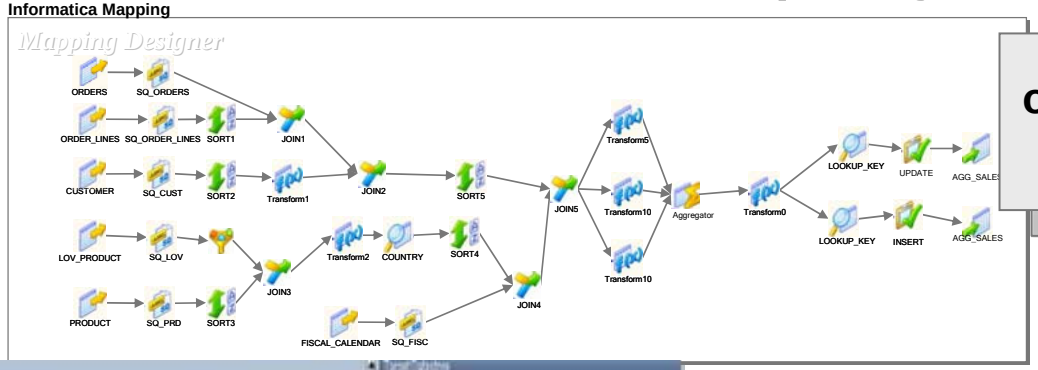

One or a related group of flow-based procedural ETL Mappings – first sample

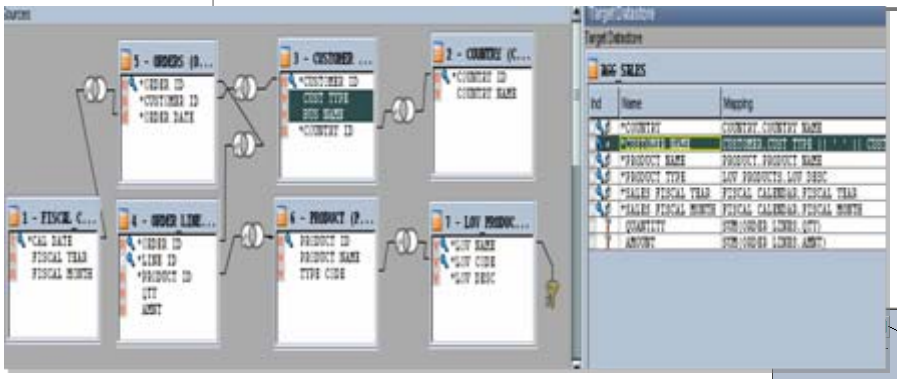

One declarative ODI interface plus selection among existing Knowledge Modules

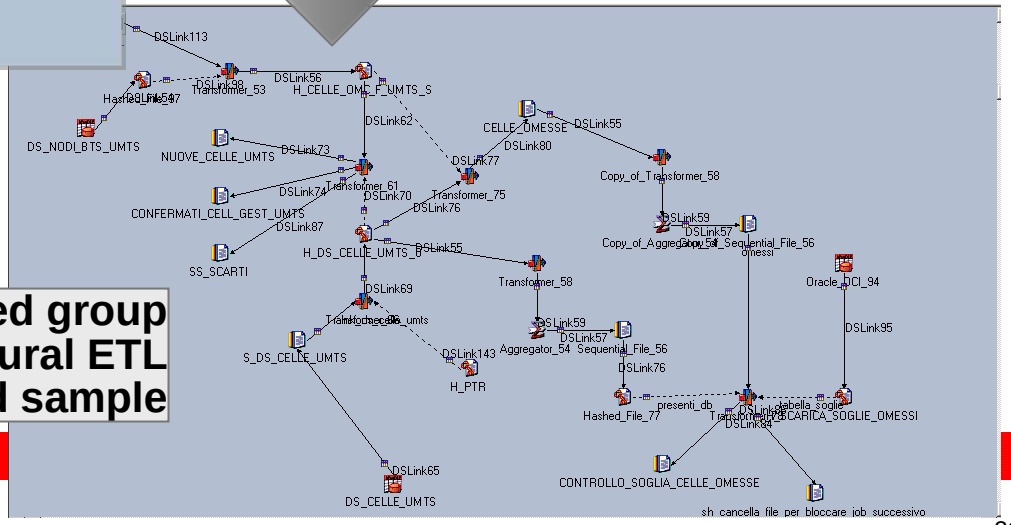

One or a related group of flow-based procedural ETL Mappings - second sample

Traditional procedural ETL

Traditional ETL row to row complexity

Informatica Mapping

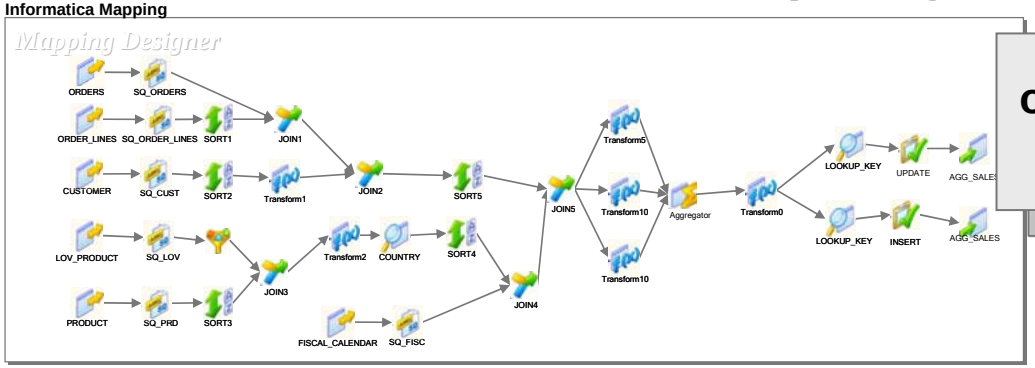

One or a related group of flow-based procedural ETL Mappings – first sample

Flow Generation is AUTOMATIC, written by ODI directly!

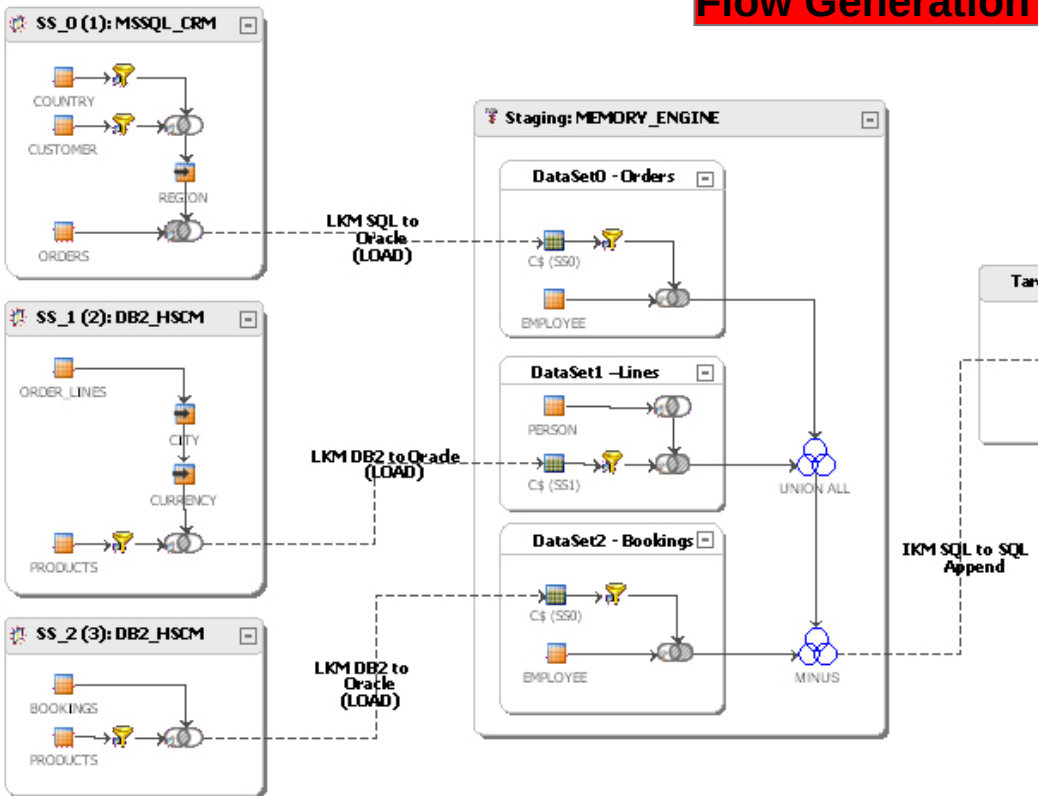

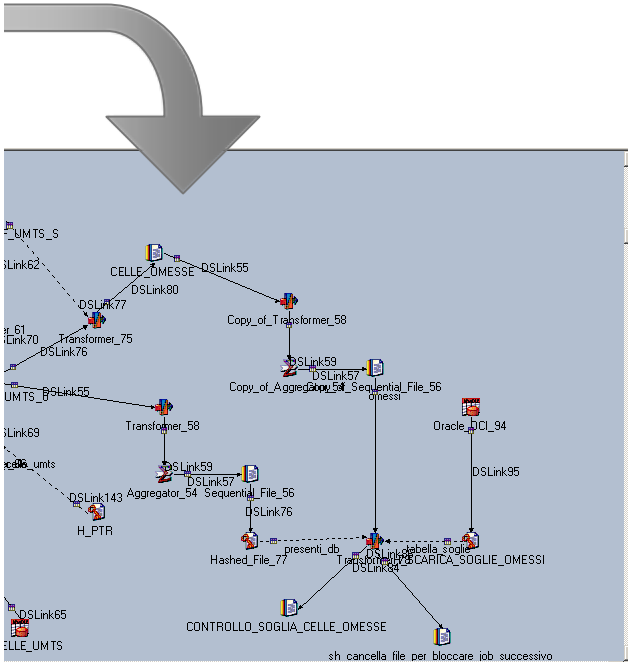

Oracle Data Quality

Ensure Data Quality as Part of the Integration Process

Visual Data Quality Tools

ODI EE

- Better Data Visibility
- Tighter Data Control
- Greater Data Accuracy

Oracle Data Profiling

Metadata Profiling of Source & Target

Oracle Data Quality

Duplicate Detection, Matching & Merging

Real Time Business Intelligence

Active Data Warehouse Improves Customer Service

Business Challenges:

- Provide up-to-second data in centralized enterprise data warehouse for business users (customer claims processing for major incidents)
- Integrate data from multiple operational systems (DB2 on z/OS and LUW)
- Leverage existing ETL investment to achieve near real-time data feeds.

GoldenGate Solution:

- Real-Time Data Warehousing captures ~40 million transactions per day from heterogeneous operational sources.
- Real-Time CDC co-exists with ELT for data requiring heavy transformations.

Results:

- Reduced claims litigation costs by 20% - enabled by faster and more responsive employee intelligence on claims data.
- Improved resource utilization efficiency by 5%
- Management team can now dynamically monitor reserve threshold to increase profitability.

Active Enterprise Intelligence

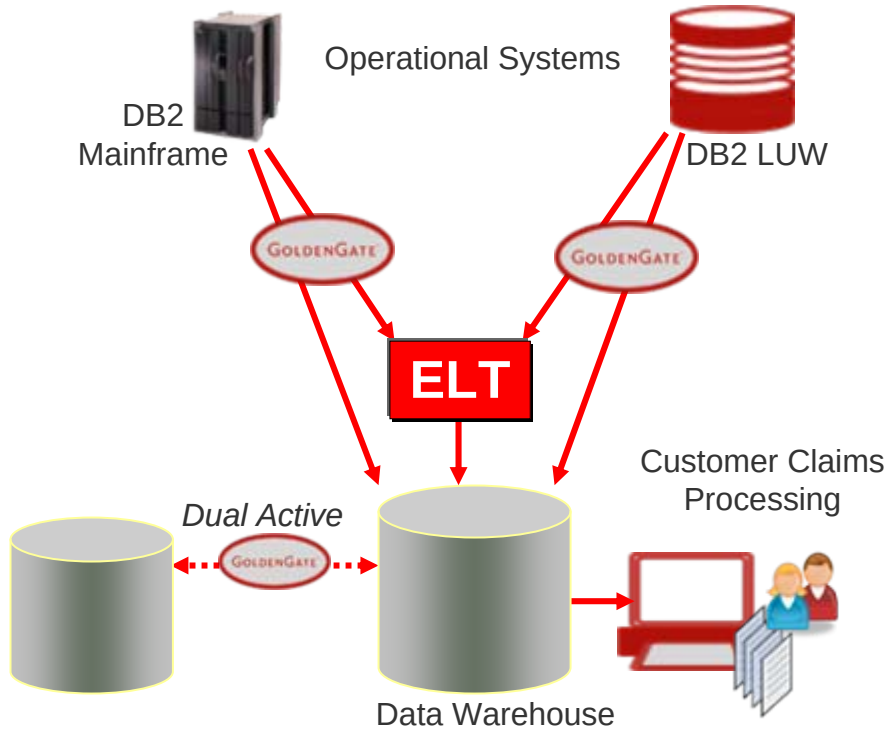

"Data is not a business problem! It is the foundation of our ability to be profitable in a rapidly changing marketplace."

- Eric Post, Director of Data Warehousing, Liberty Mutual

Customer Example: Overstock.com

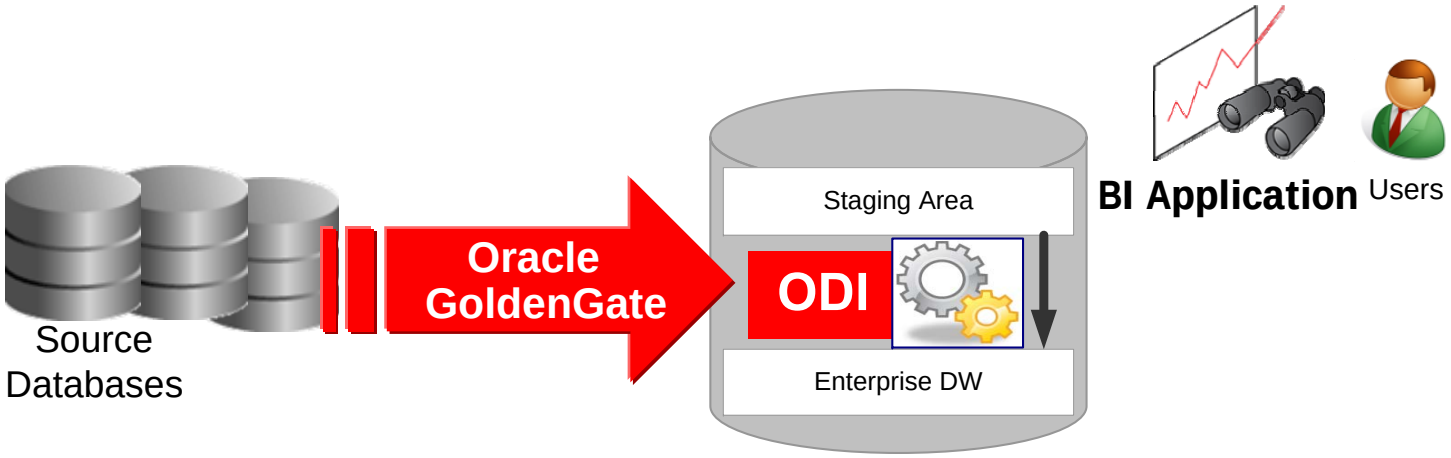

Results:

- Optimize decisions and decrease costs with real-time information
  - Customer analysis now done in minutes, rather than days
  - Reduced marketing spend with ‘smart’ campaign decisions
  - Reduced inventory costs
- Eliminated batch windows
- Ranked #4 in customer service by *National Retail Federation*

Real Time Data Integration Case Study

Better BI to Reduce Customer Churn

COMPANY OVERVIEW

DIRECTV is a $17 billion provider of satellite-based television services. DIRECTV's 7,500 employees operate the company's broadcast centers, monitor satellites, and deliver service to about 17 million U.S. and over 5 million Latin American customers.

CHALLENGES / OPPORTUNITIES

- **Maintain high quality customer service in competitive market – reduce churn!**
- **Centralize customer information for a single view to support sales, marketing, support & field service**
- **Significantly reduce data latency in central data warehouse for all queries & reports – edict for < 15 minutes!**

GoldenGate PROVIDES

- **Real-time data integration from Siebel CRM on Oracle to central Warehouse**
- **1,500 service agents log 600,000 customer calls p/day**
- **GoldenGate moves 150-200 million records per day with 1.5 second latency.**

RESULTS

- **Significantly reduced churn by 25%**
- **All business units have access to real-time business data.**

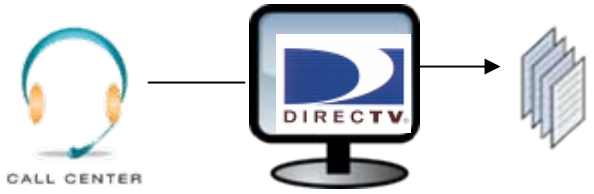

High Availability: Live Standby

Software as a Service to Financial Institutions, No Interruption

Business Challenges:

- Ensure no interruption to ongoing business operations for banking and financial institution customers
- Enable continuous availability of critical applications—account opening, funds transfer
- Implement a reliable, in-house live standby solution

GoldenGate Solution:

- GoldenGate Software Live Standby for High Availability / Disaster Tolerance replicates data from multiple Oracle RAC clusters to an exact replica of the production environment
- In case of outage or failure, web applications are unaffected and "business as usual" is maintained for global financial services clients
- Live Standby increases ROI by allowing for data integration into data warehouse to run reports via Oracle BI Enterprise Edition

Live Standby for High Availability and Disaster Tolerance

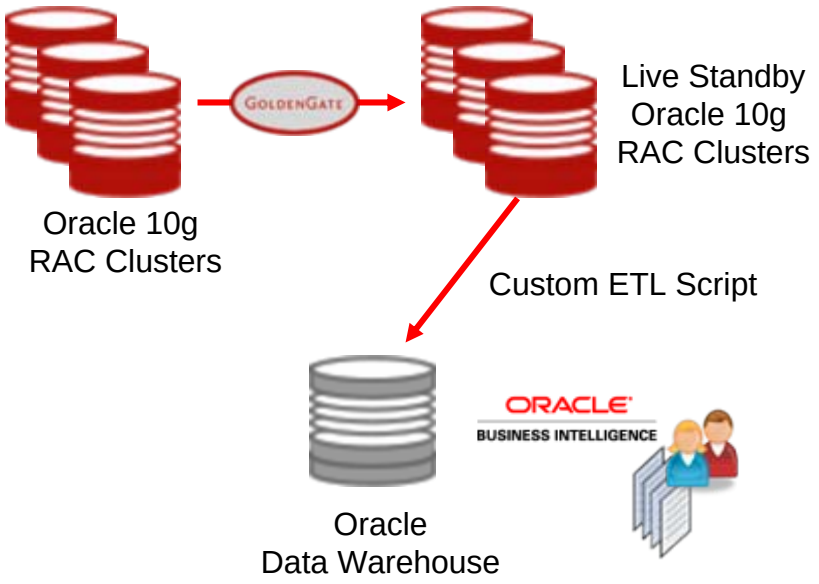

"GoldenGate has opened our eyes to other areas throughout our business where near real-time data can deliver added value to the business and a better return on investment."
- Burak Yenier, Sr. Dir. of Operations, CashEdge

Oracle Data Integration platform Use Cases

Comprehensive Best-of-breed Approach for Unifying Information

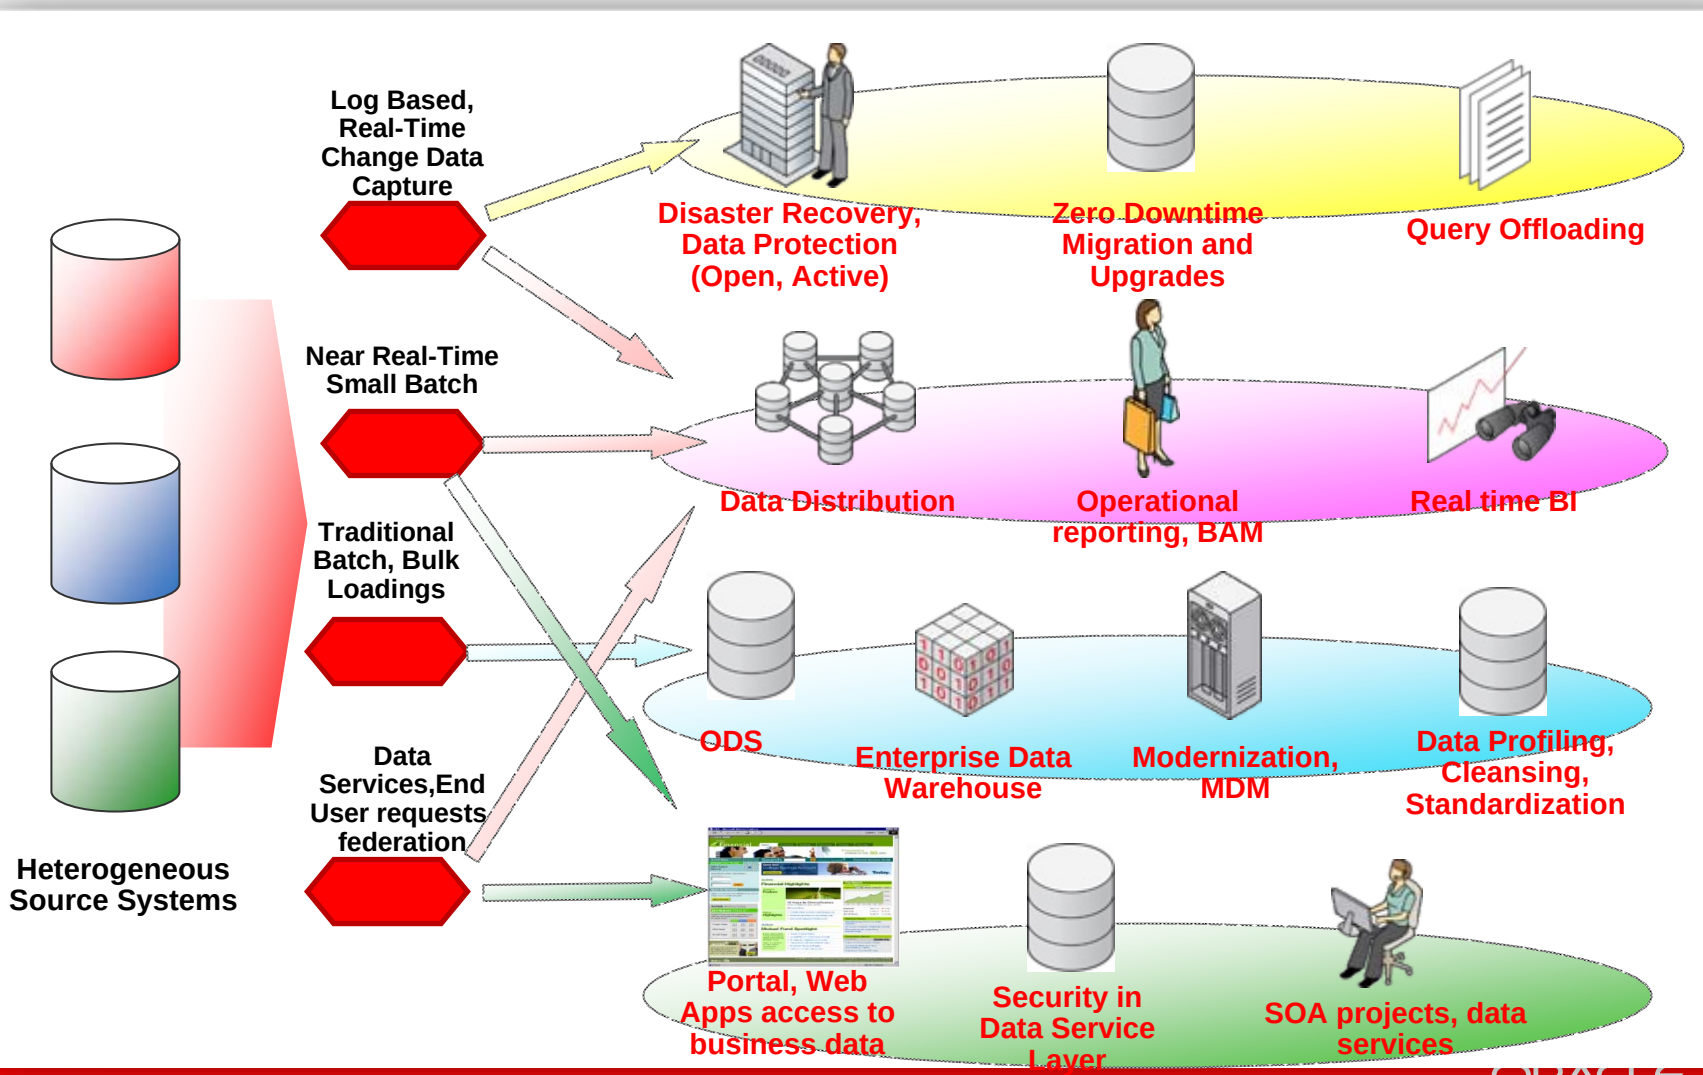


Q&A

The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle’s products remains at the sole discretion of Oracle.


ORACLE®