ORACLE

Introducing Oracle Data Integrator and Oracle GoldenGate

Marco Ragogna
EMEA Principal Sales Consultant
Data integration Solutions

IT Obstacles to Unifying Information

What is it costing you to unify your data?

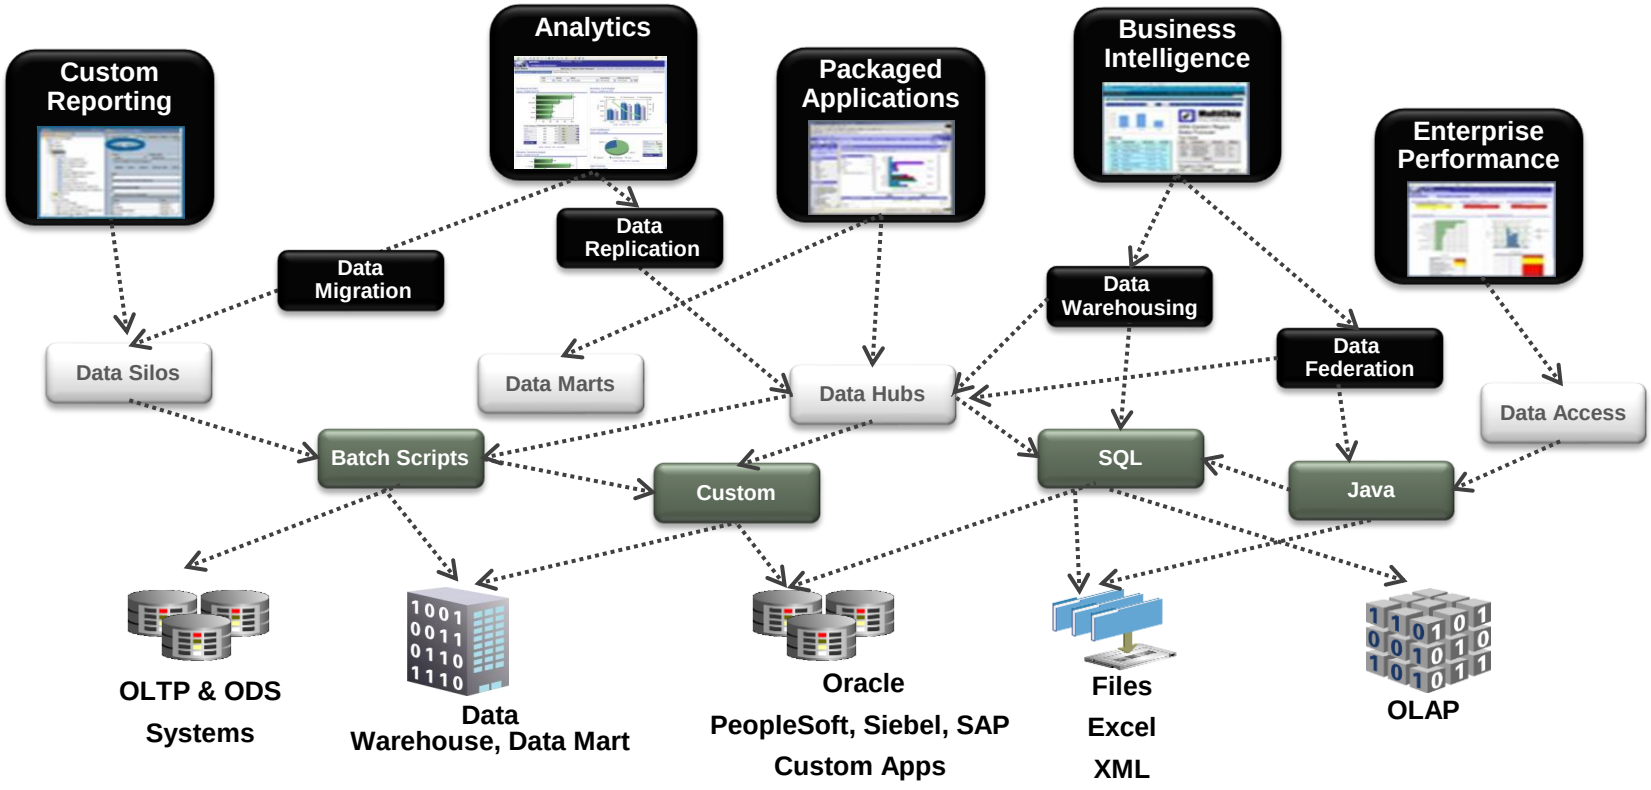

Fragmented Data Silos | **Slow Performance** | **Out of sync** | **Poor Data Quality** | **What's the cost?**

Data Integration

Key Component of Oracle Fusion Middleware

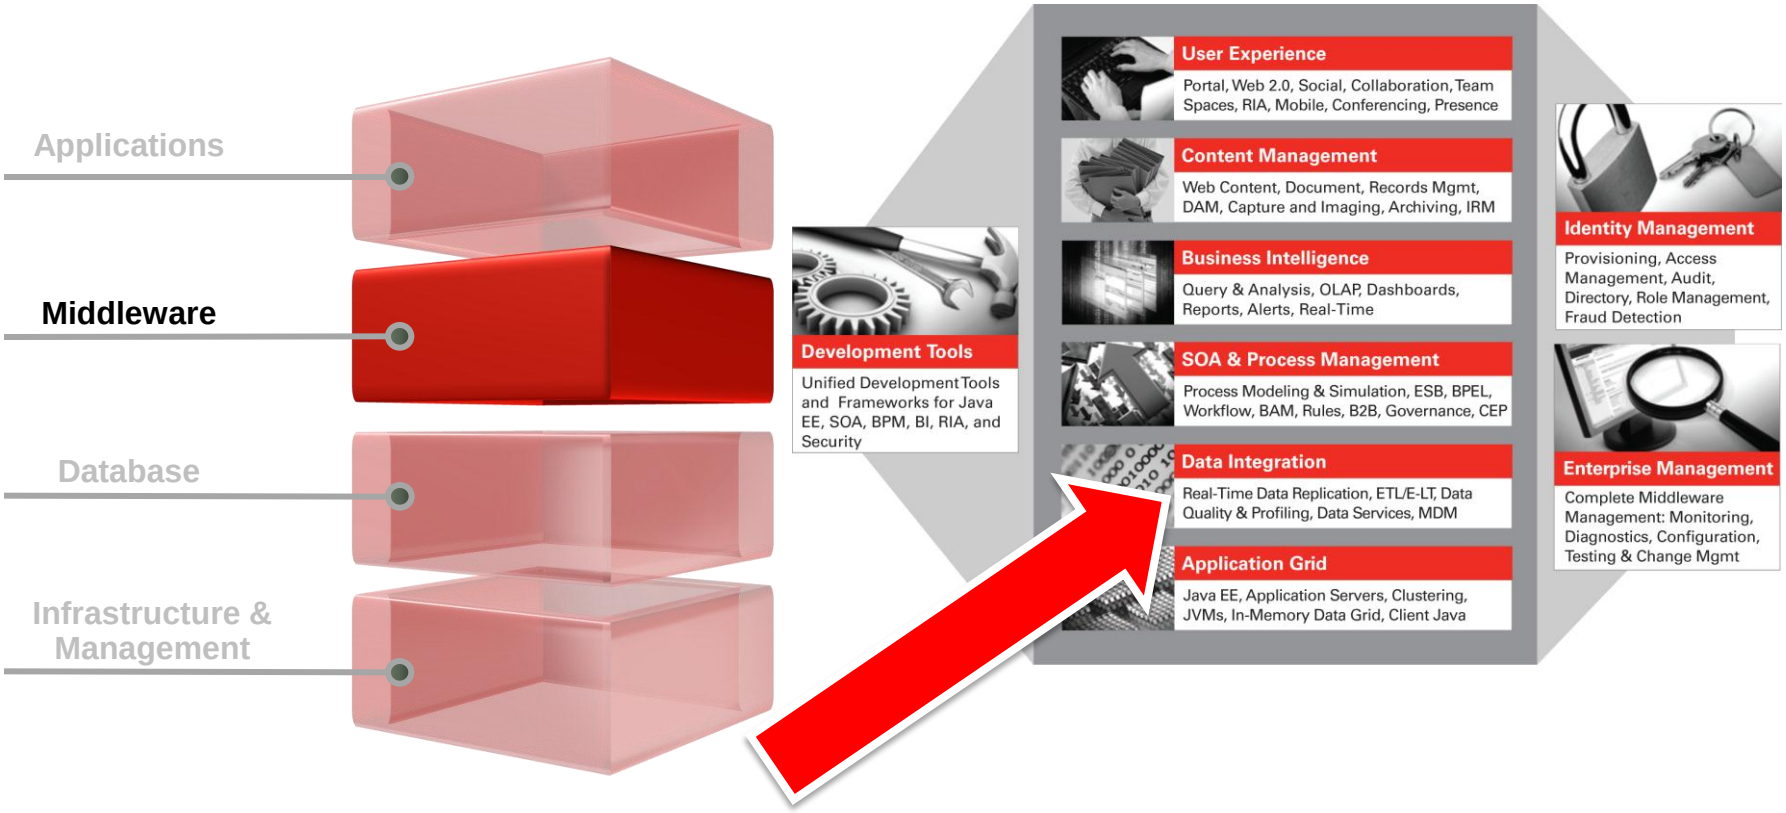

Oracle Data Integration

The solution for enterprise-wide real-time data

Databases

Distributed systems

Legacy systems

OLAP systems

OLTP systems

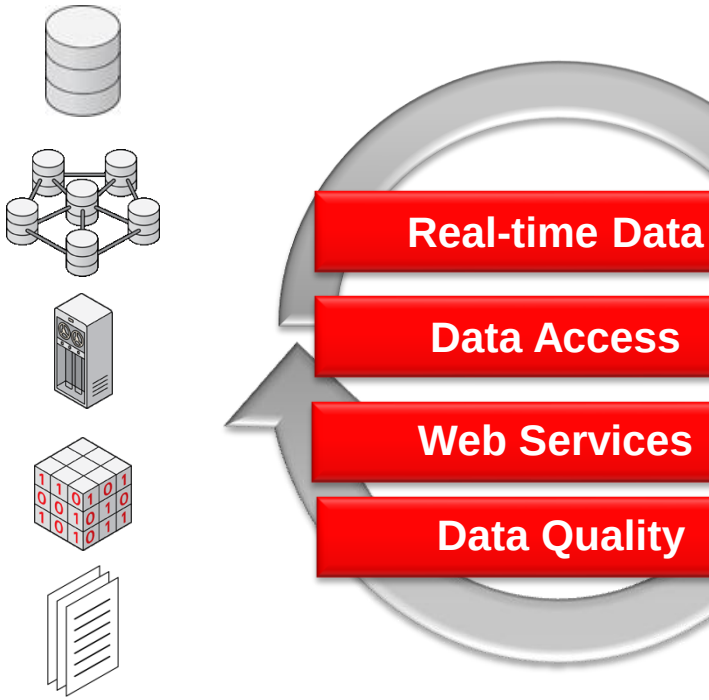

Mission critical systems and data

Business Intelligence, Performance Management

Data Warehouses, MDM

SOA

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

Oracle Data Integration

The solution for enterprise-wide real-time data

Databases

Distributed systems

Legacy systems

OLAP systems

OLTP systems

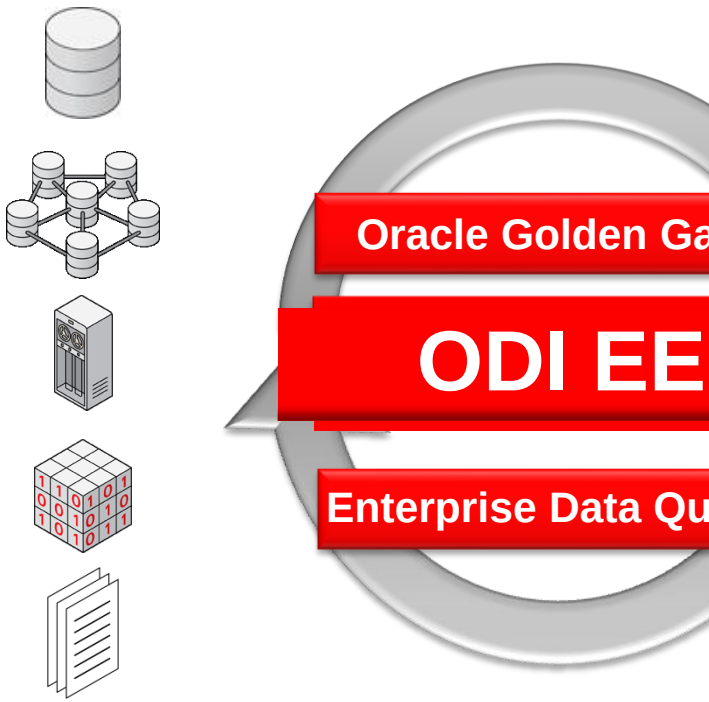

Mission critical systems and data

Business Intelligence, Performance Management

Data Warehouses, MDM

SOA

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

Oracle Data Integration

The solution for enterprise-wide real-time data

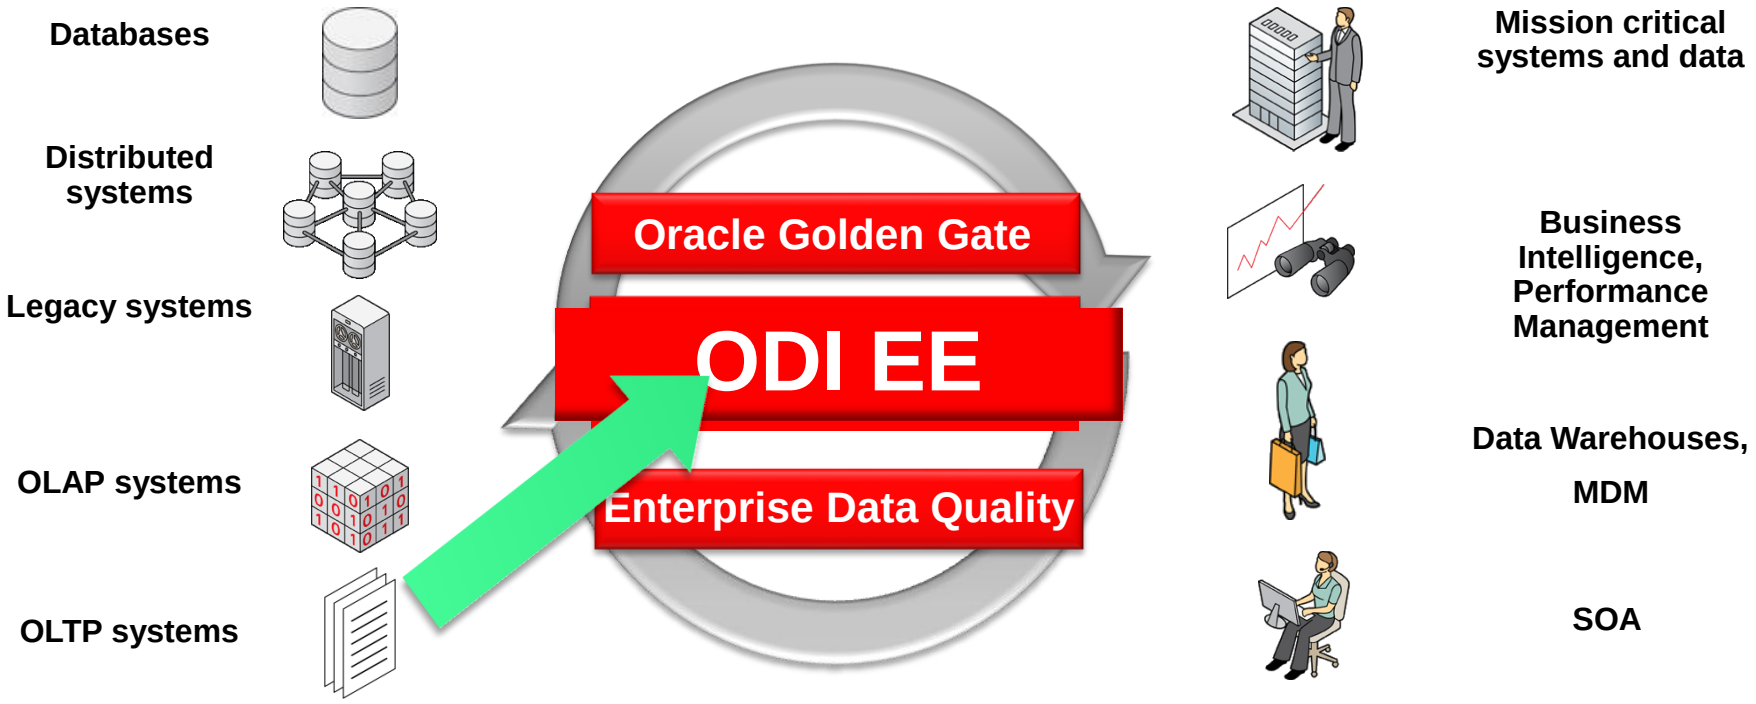

Dramatically improve the accessibility, reliability, and quality of critical data across enterprise systems

Why Does ODI Win?

ODI is Faster

- Fastest E-LT Bulk/Batch Performance
- Faster Real Time integration (sub-second trickle) with CDC, Replication, and SOA infrastructure
- Faster Project Setup, Design and Delivery

ODI is Simpler

- Simpler Setup, Configuration, Management, and Monitoring
- Simpler way to do Mapping using Declarative SQL Interfaces
- Simpler Deployment with Fewer Hardware Devices
- Simpler extensibility with Knowledge Module code templates

ODI is Saves Money (Lower TCO, Higher ROI)

- Less Hardware & Energy Costs with E-LT Architecture
- Less Time Wasted on Unnecessary ETL Mappings, Scripting, and Complex Training
- Less Integration Overhead Integrating with Applications, SOA, and Management Software

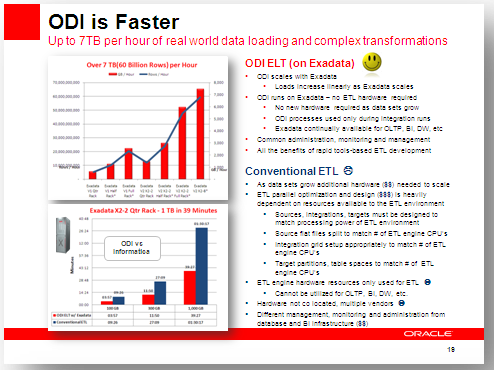

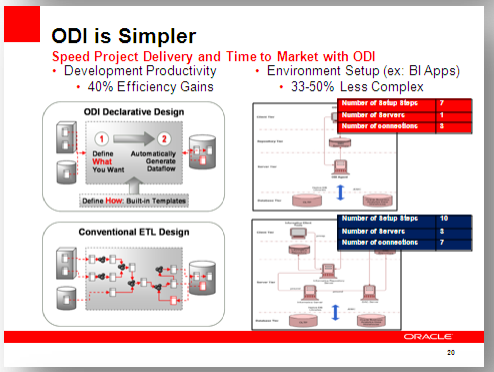

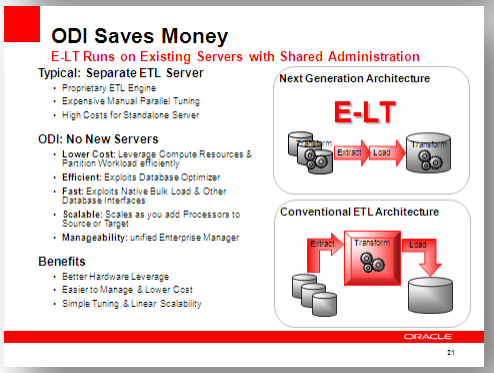

ODI Saves Money

E-LT Runs on Existing Servers with Shared Administration

Typical: Separate ETL Server

- Proprietary ETL Engine
- Expensive Manual Parallel Tuning
- High Costs for Standalone Server

ODI: No New Servers

- **Lower Cost:** Leverage Compute Resources & Partition Workload efficiently
- **Efficient:** Exploits Database Optimizer
- **Fast:** Exploits Native Bulk Load & Other Database Interfaces
- **Scalable:** Scales as you add Processors to Source or Target
- **Manageability:** unified Enterprise Manager

Benefits

- Better Hardware Leverage
- Easier to Manage & Lower Cost
- Simple Tuning & Linear Scalability

Next Generation Architecture

E-LT

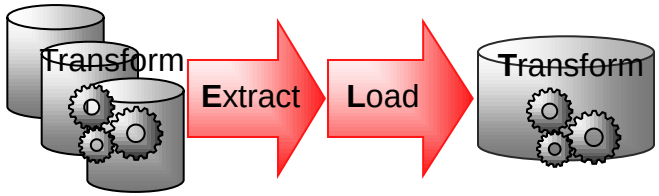

Conventional ETL Architecture

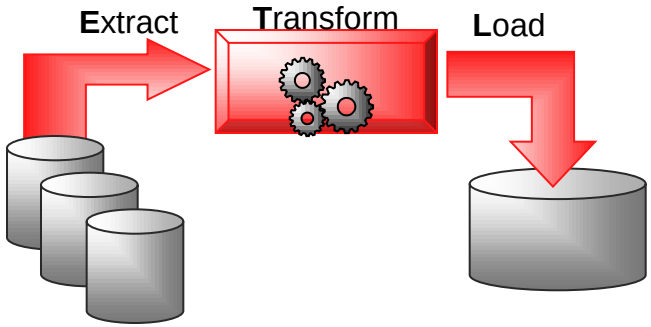

ODI is Faster

Up to 7TB per hour of real world data loading and complex transformations

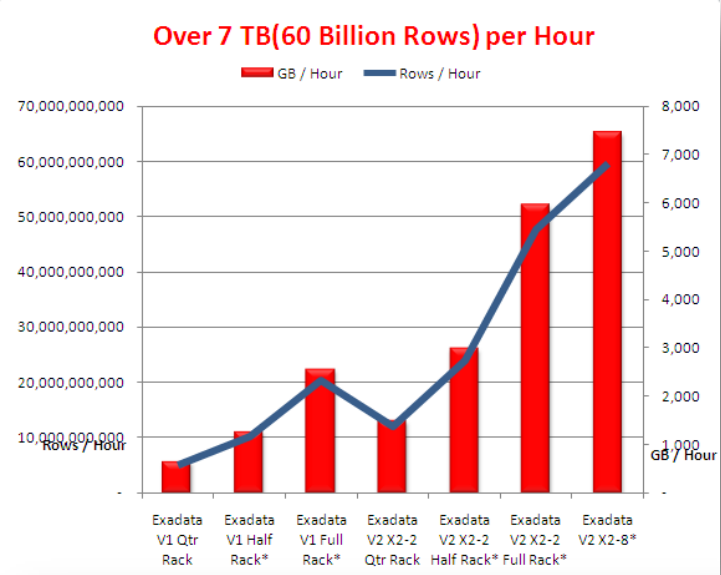

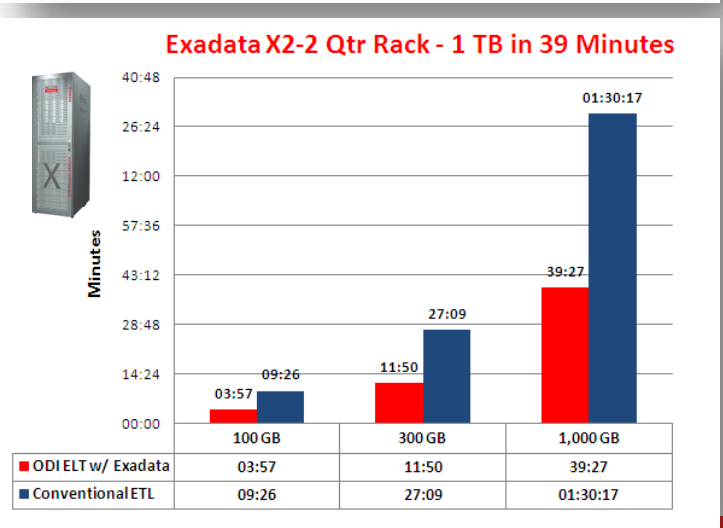

ODI ELT (on Exadata/any DW)

- ODI scales with the Database
  - Loads increase linearly as DW scales
- ODI runs on relational technologies – no ETL hardware required
  - No new hardware required as data sets grow
  - ODI processes used only during integration runs
  - Databases continually available for OLTP, BI, DW, etc
- Common administration, monitoring and management
- All the benefits of rapid tools-based ETL development

Conventional ETL

- As data sets grow, more hardware (**$$**) needed to scale
- ETL parallel optimization and design (**$$$**) is heavily dependent on resources available to the ETL environment
  - Sources, integrations, targets must be designed to match processing power of ETL environment
  - Source flat files split to match # of ETL engine CPU's
  - Integration grid setup appropriately to match # of ETL engine CPU's
  - Target partitions, table spaces to match # of ETL engine CPU's
- ETL engine hardware resources only used for ETL
  - Cannot be utilized for OLTP, BI, DW, etc.
- Hardware not co located, multiple vendors
- Different management, monitoring and administration from database and BI infrastructure (**$$**)

“Old Style” ETL

- Monolithic & Expensive Environments
- Fragile, Hard to Manage
- Difficult to Tune or Optimize

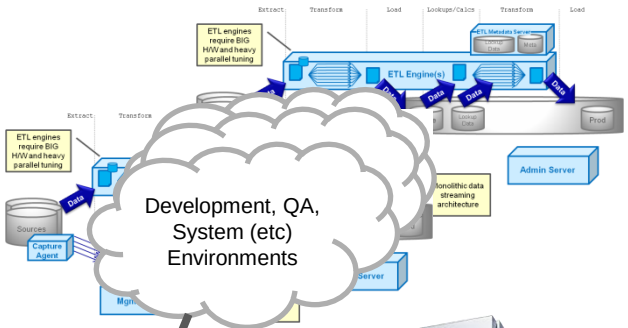

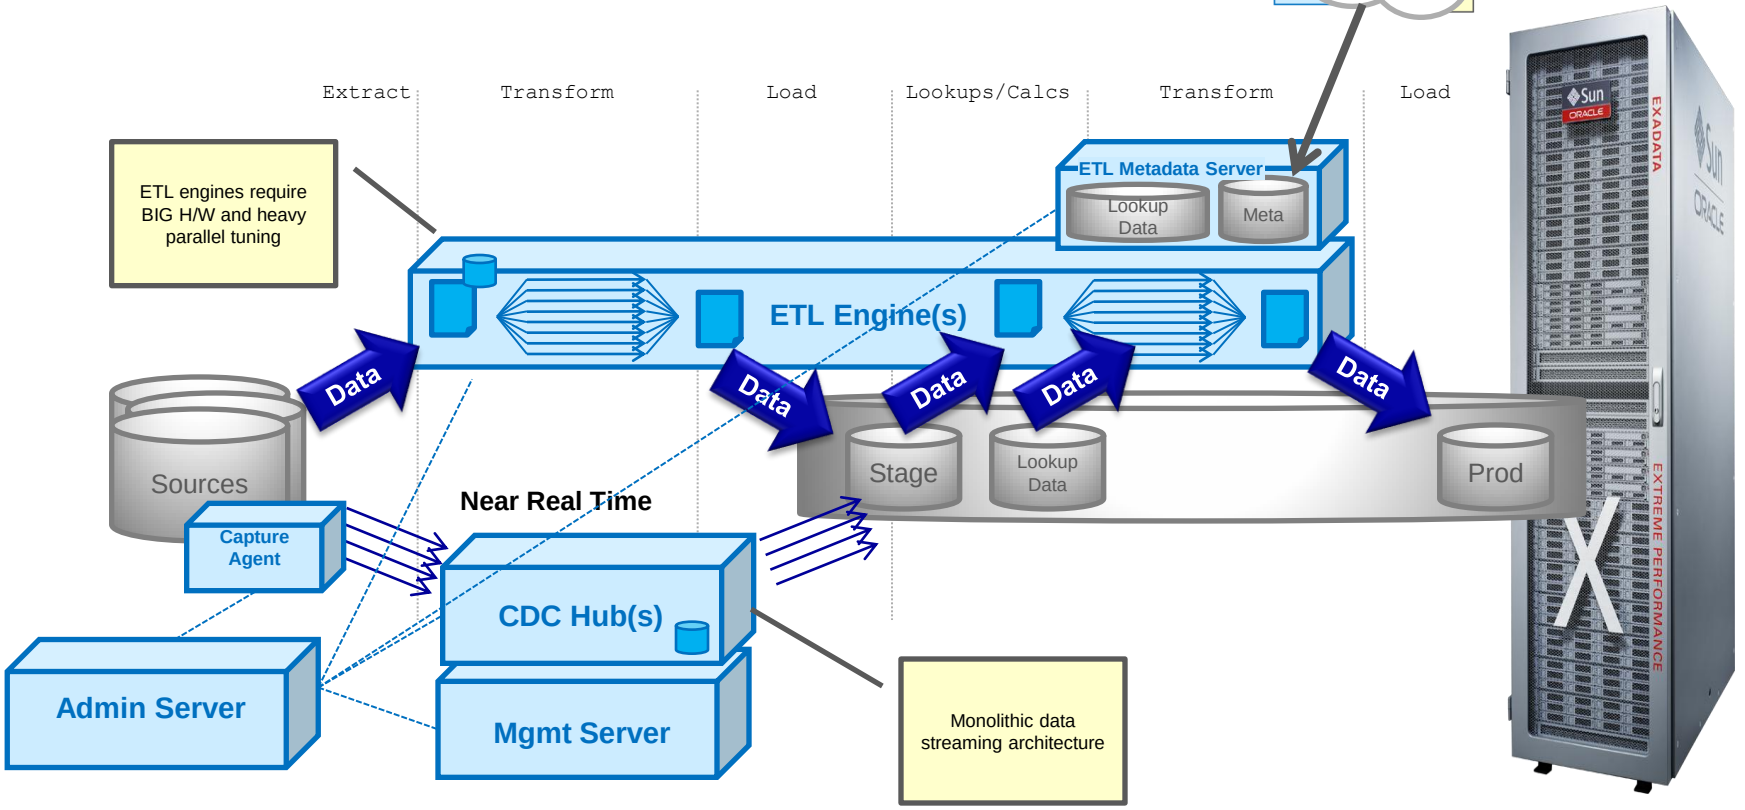

Modern Data Integration

- Lightweight, Inexpensive Environments – Agents
- Resilient, Easy to Manage – Non-Invasive
- Easy to Optimize and Tune – uses DBMS power

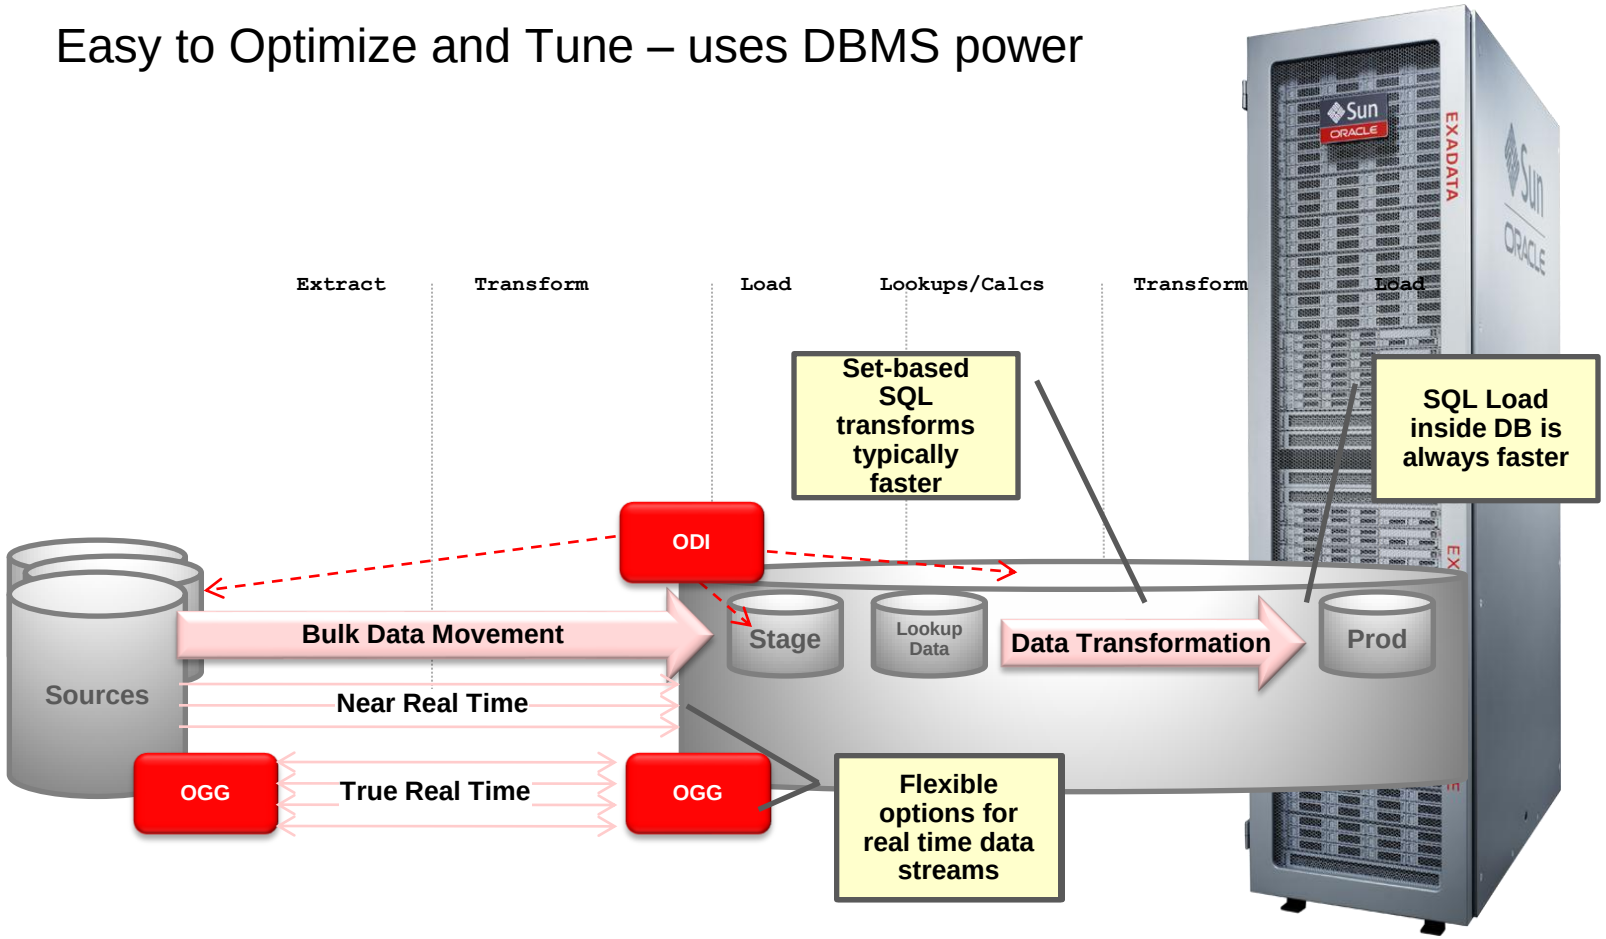

Best Data Integration for Exadata

Top Performance, Smallest Footprint

- Run ODI, EDQ & OGG Directly on Exadata
- Support Any Latency Data Feeds
- Non-Invasive Source Capture
- Most Cost-Effective and High-Performance Exadata Data Loading

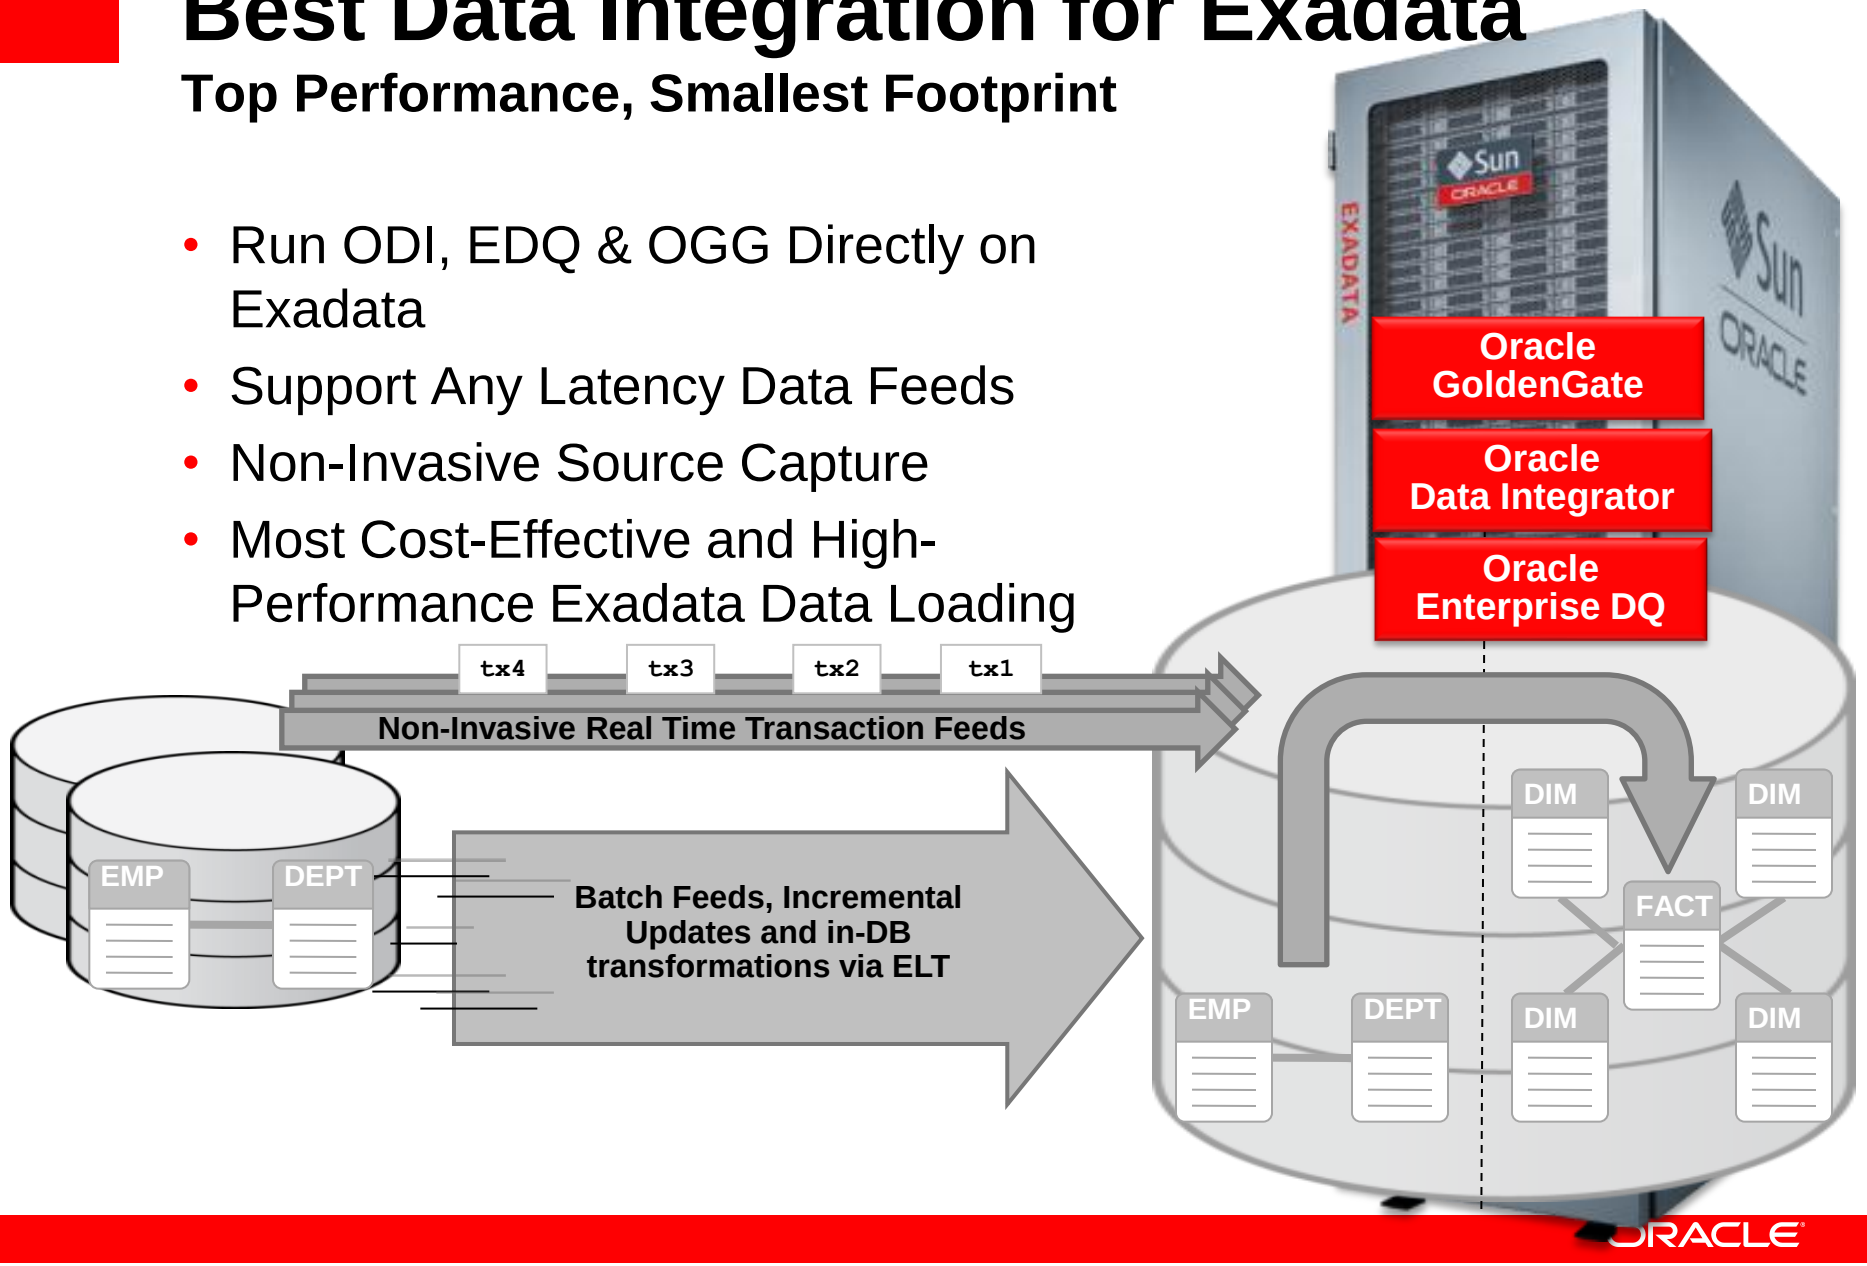

ODI is Simpler

Speed Project Delivery and Time to Market with ODI

- Development Productivity
  - 40% Efficiency Gains
- Environment Setup (ex: BI Apps)
  - 33-50% Less Complex

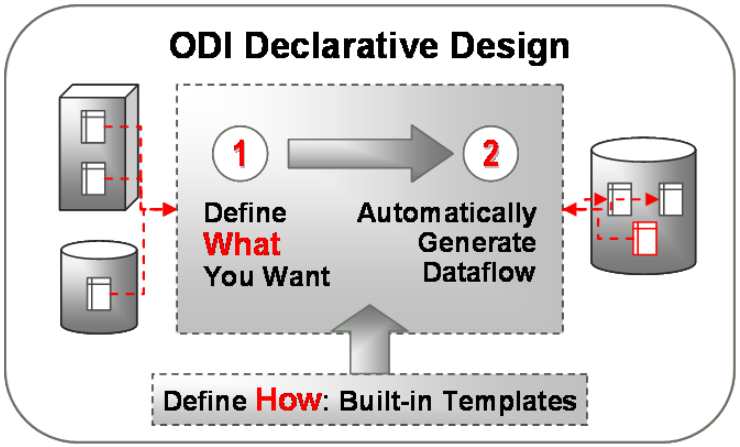

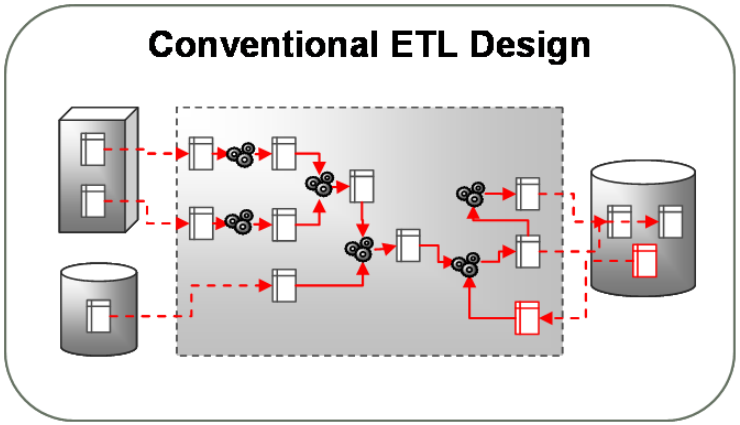

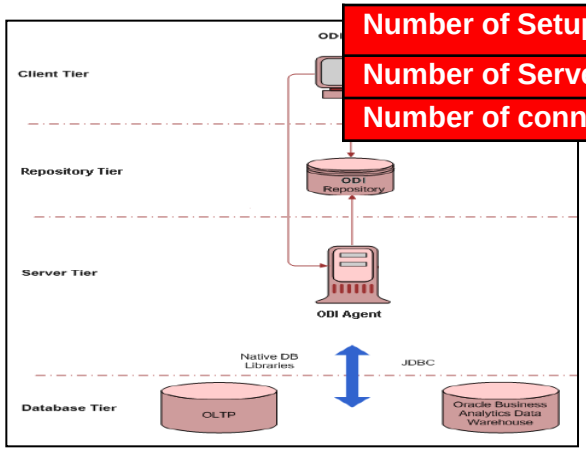

Number of Setup Steps	7
Number of Servers	1
Number of connections	3

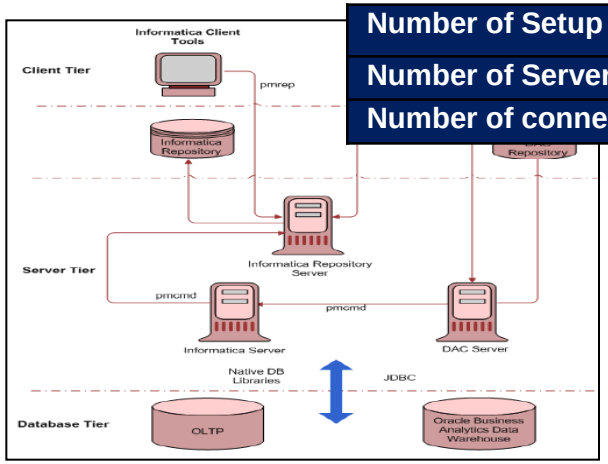

Number of Setup Steps	10
Number of Servers	3
Number of connections	7

Traditional procedural ETL

Traditional ETL row to row complexity

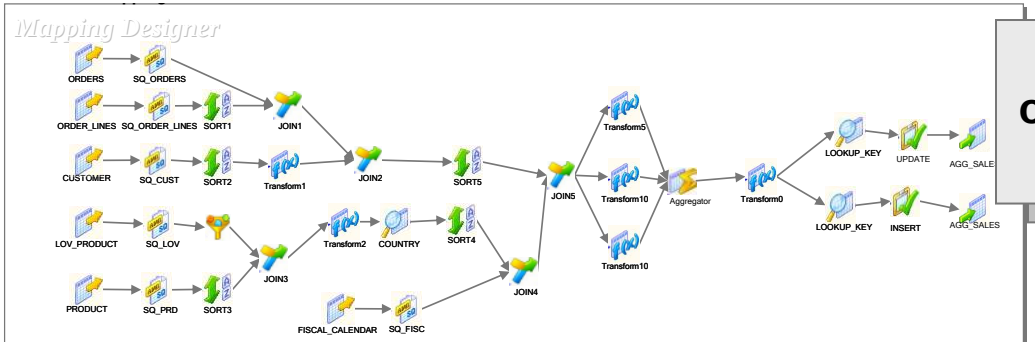

One or a related group of flow-based procedural ETL Mappings – first sample

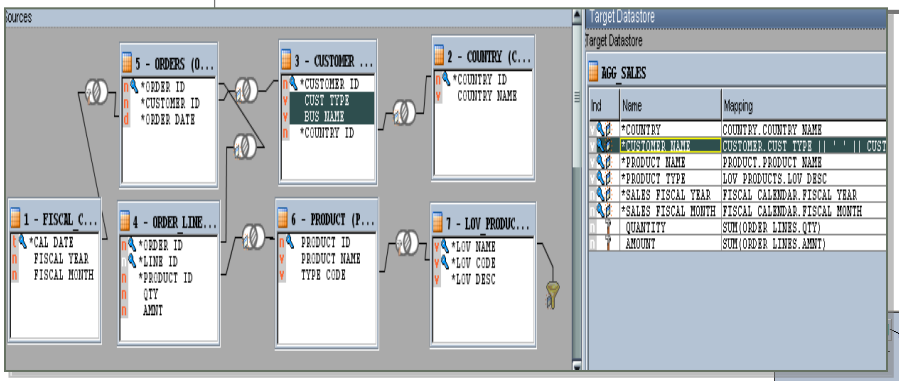

One declarative ODI interface plus selection among existing Knowledge Modules

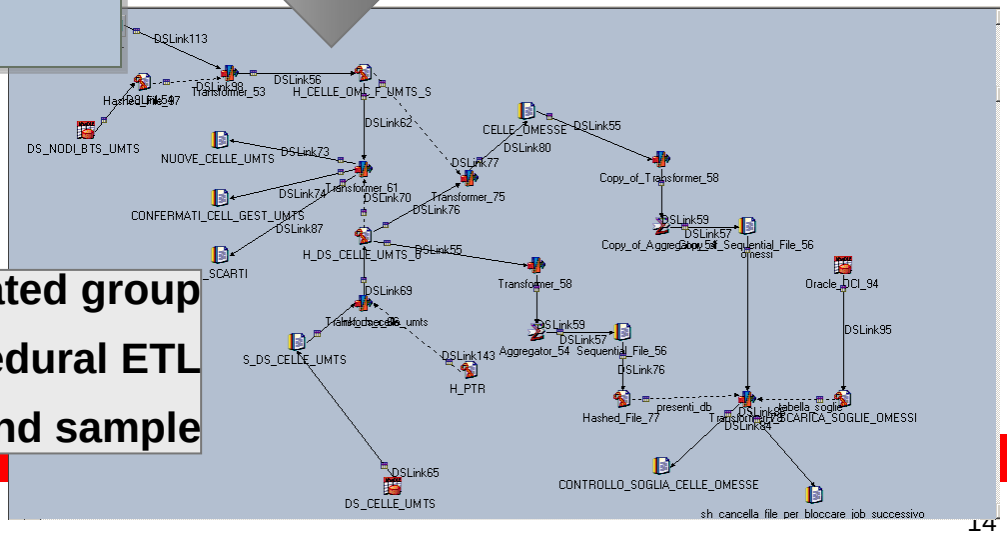

One or a related group of flow-based procedural ETL Mappings - second sample

Traditional procedural ETL

Traditional ETL row to row complexity

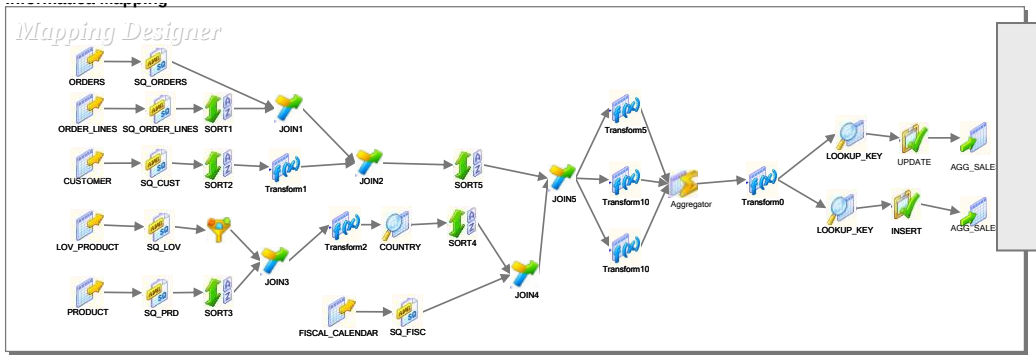

One or a related group of flow-based procedural ETL Mappings – first sample

Flow Generation is AUTOMATIC, written by ODI directly!

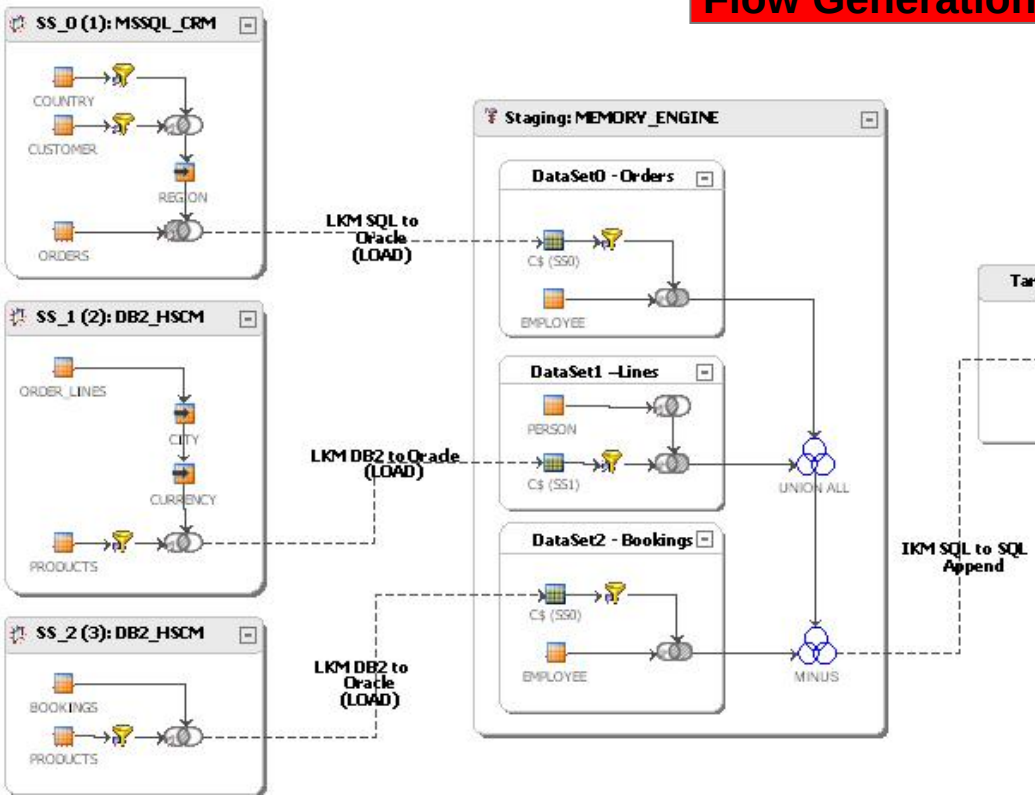

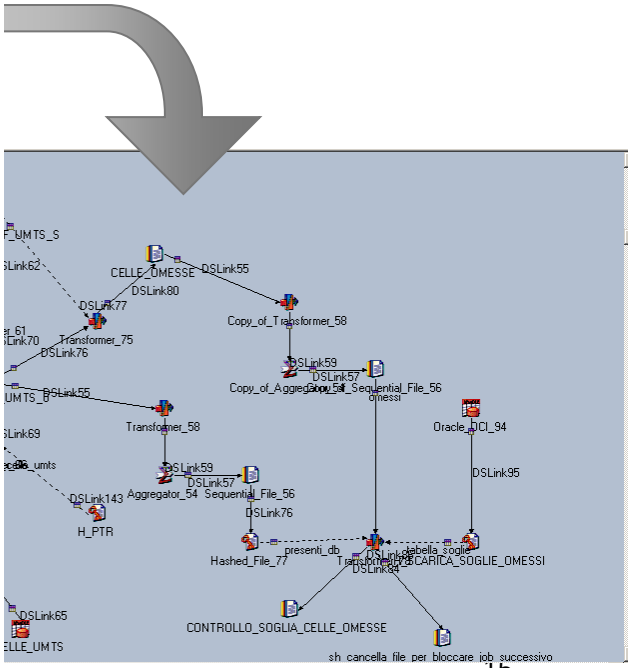

Topology Module on ODI

- You describe how the relational infrastructure where ODI works is done
  - ODI builds the flow for a specific loading automatically!

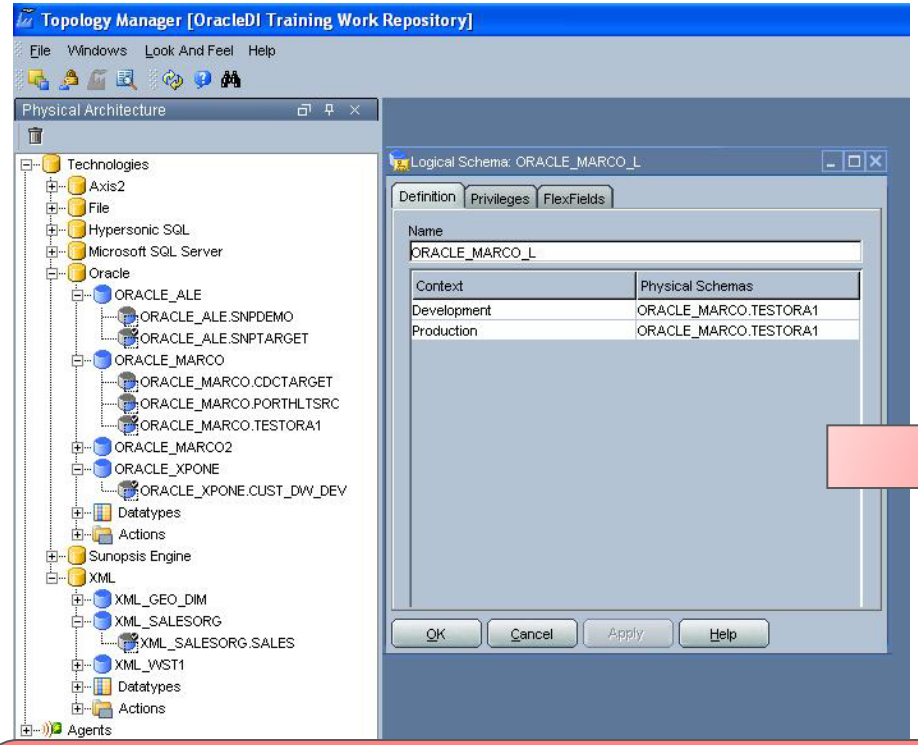

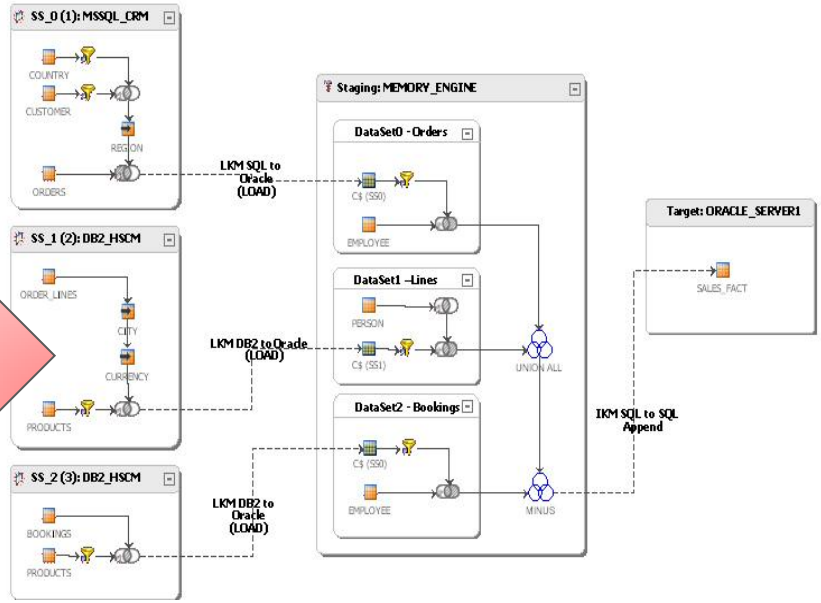

Topology module allows to describe all the information on the technology where the ELT projects work, starting from specific definition on the technologies that are used, going to physical description on how to access a server, wich user and password to enter, which schema users or database are involved in the jobs. The final developer will have only a logical reference to the servers

Declarative mapping + Knowledge Modules = Generated Code

KM's Meta Code

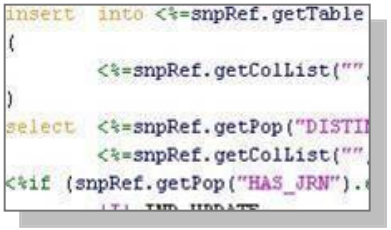

Metadata

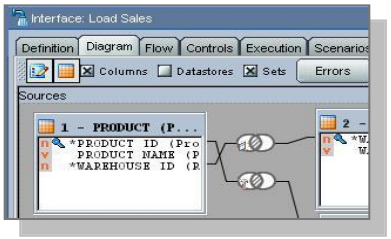

KM Interpreter

Executed Code

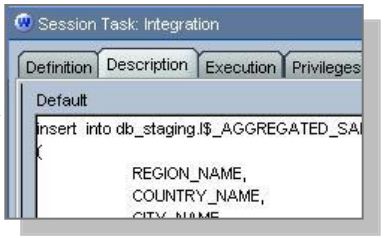

- 120+ KMs out-of-the-box
  - ✓ Tailor to existing best practices
  - ✓ Ease administration work
  - ✓ Reduce cost of ownership
- Customizable and extensible

Siebel
DB2 Exp/Imp
Oracle SQL*Loader
JMS Queues
Type II SCD
SAP/R3
TPump/ Multiload
Oracle Merge
Oracle Web Services

Pluggable Knowledge Modules Architecture

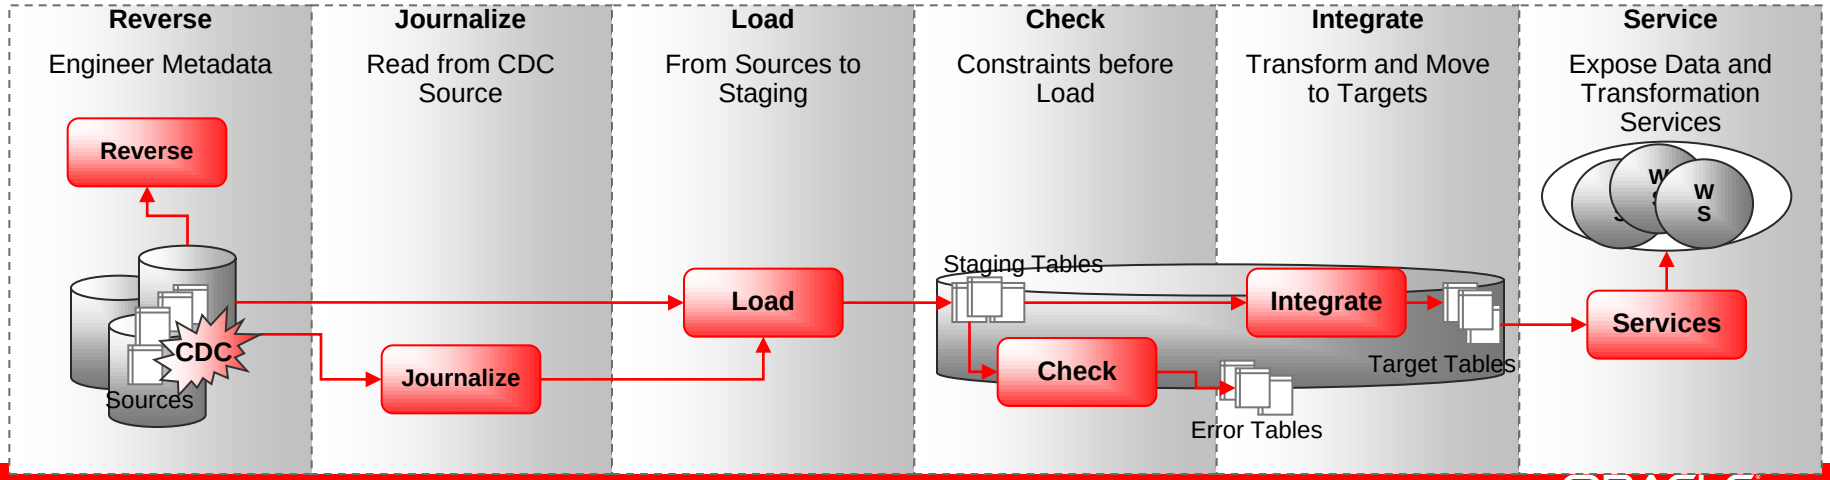

Jobs, auditing

- Technical and business metadata: ability to manage in a unique and centralized way jobs, their transformation, schedulings, data definition language etc.
  - Central Monitoring and Logging: verifying the execution of jobs

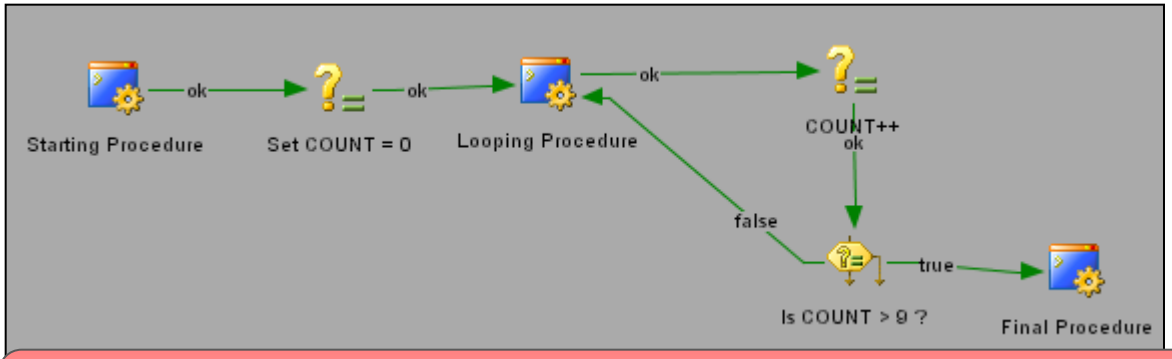

Graphical environment allows to describe job complex as needed, created putting together simple steps like the declarative design

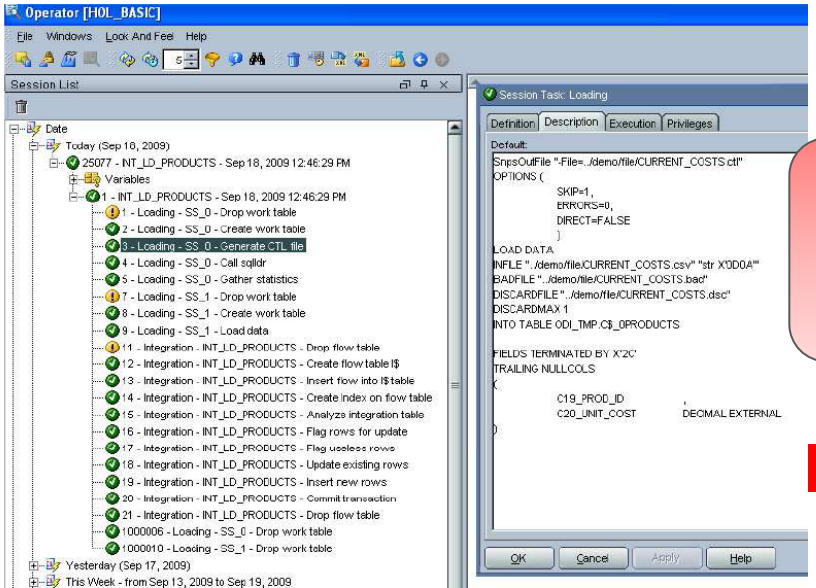

ELT Agent writes back on the repository the auditing offor the job executions, giving information on generated code, warnings and database errors that can eventually occur

Oracle Data Integration

The solution for enterprise-wide real-time data

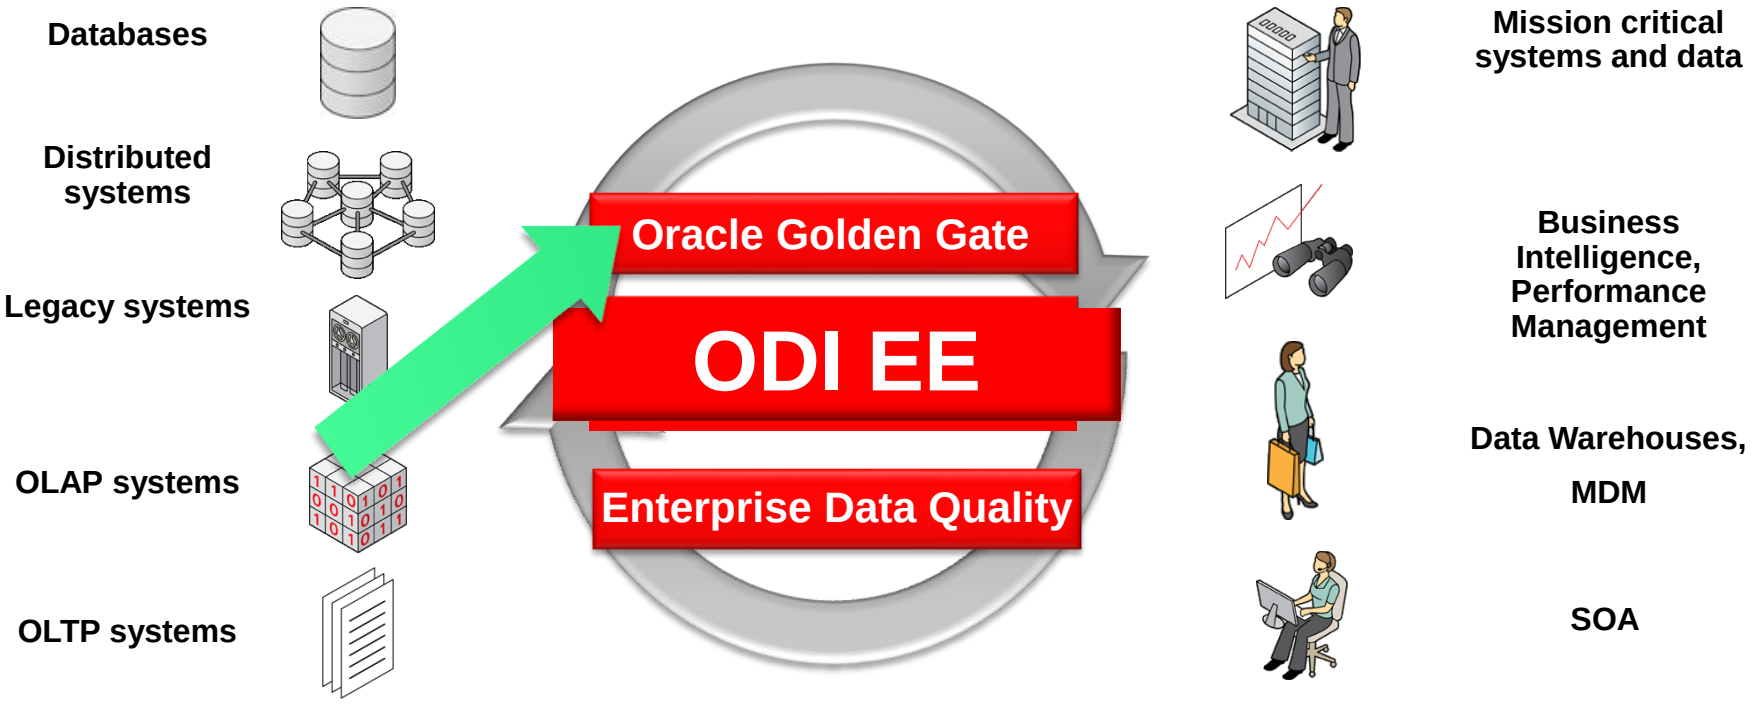

Dramatically improve the accessibility, reliability, and quality of critical data across enterprise systems

Oracle GoldenGate Overview

Oracle GoldenGate provides low-impact capture, routing, transformation, 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 GoldenGate Use Cases

Enterprise-wide Solution for Real Time Data Needs

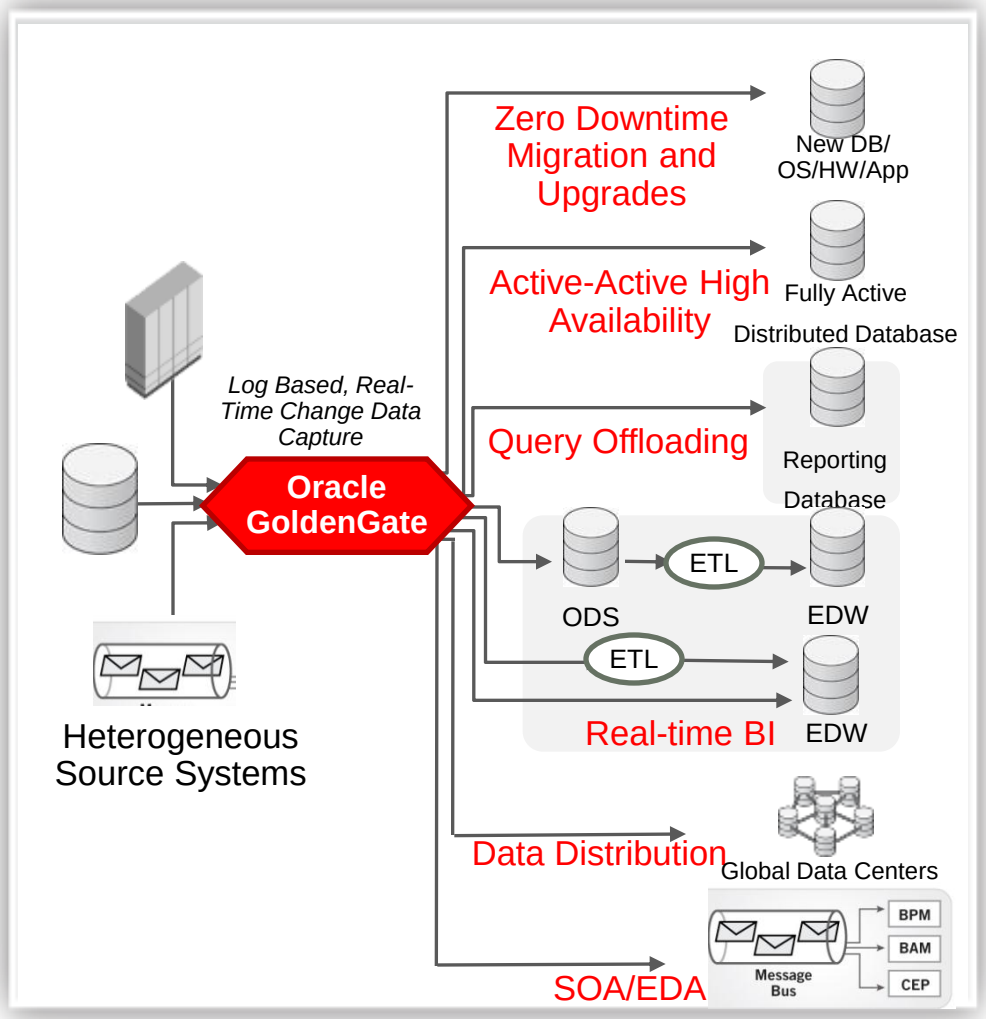

Reduce Costs

Lower Risks

Achieve Operational Excellence

Advantages of Oracle GoldenGate Architecture

Reduced Overhead and TCO

- Captures once, delivers to many targets for different uses
- Non-invasive, log-based capture
- Moves only committed data, reduces bandwidth needs

High Performance with Reliability

- Subsecond latency even with high data volumes
- Preserves transaction integrity
- Ensures data recoverability

Flexibility and Ease of Use

- Provides decoupled, modular architecture
- Supports heterogeneous sources and targets, and different latency needs
- Coexists and integrates with ELT/ETL and messaging solutions

How Oracle GoldenGate Works

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

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.

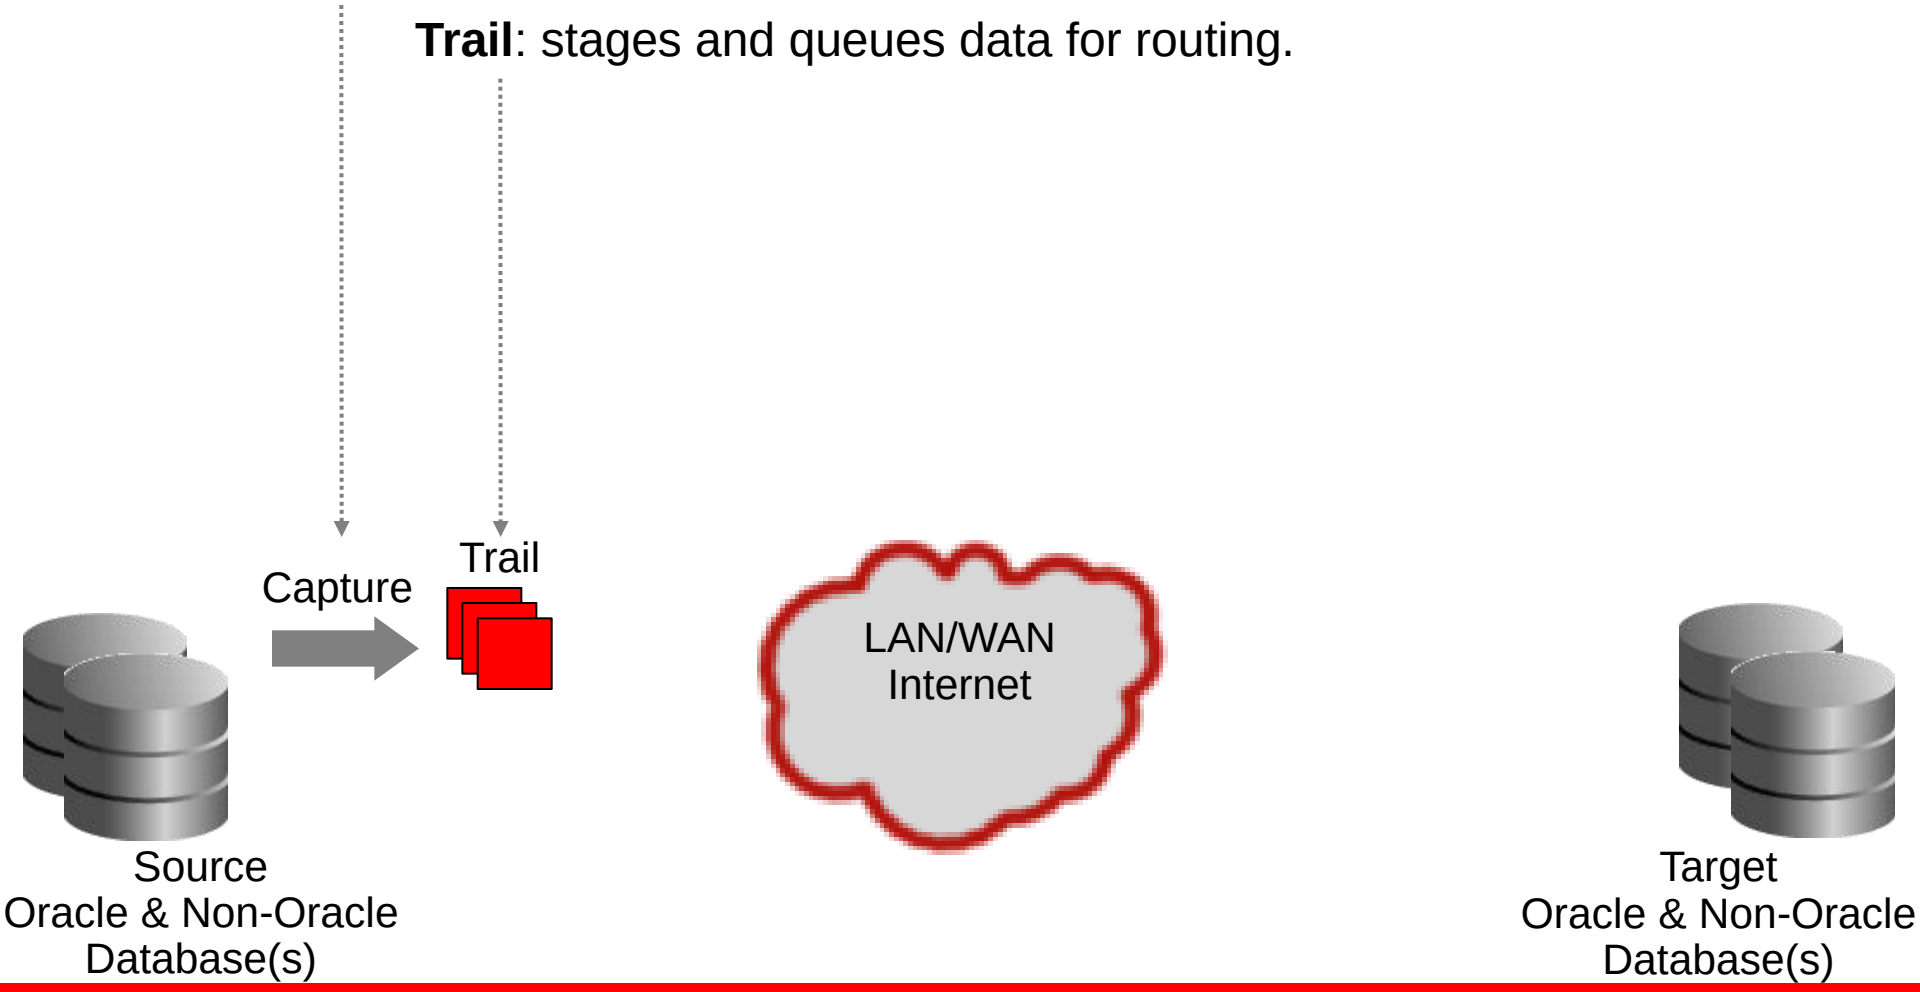

How Oracle GoldenGate Works

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).

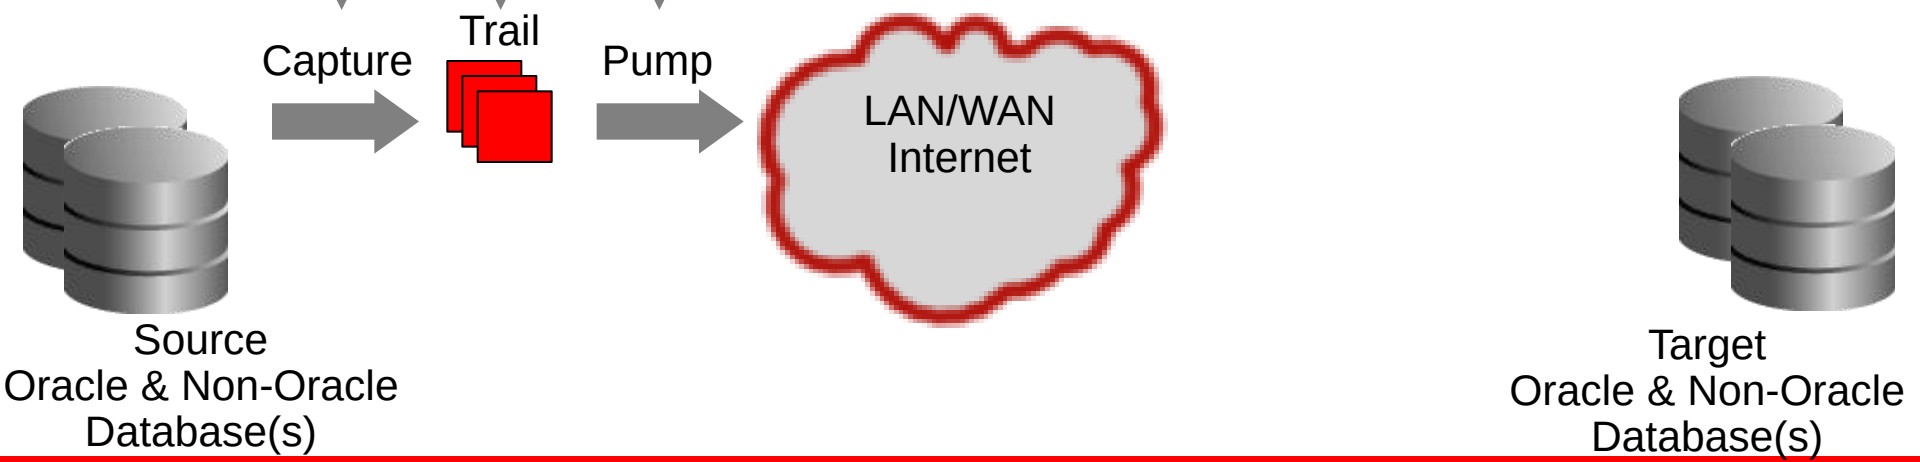

How Oracle GoldenGate Works

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).

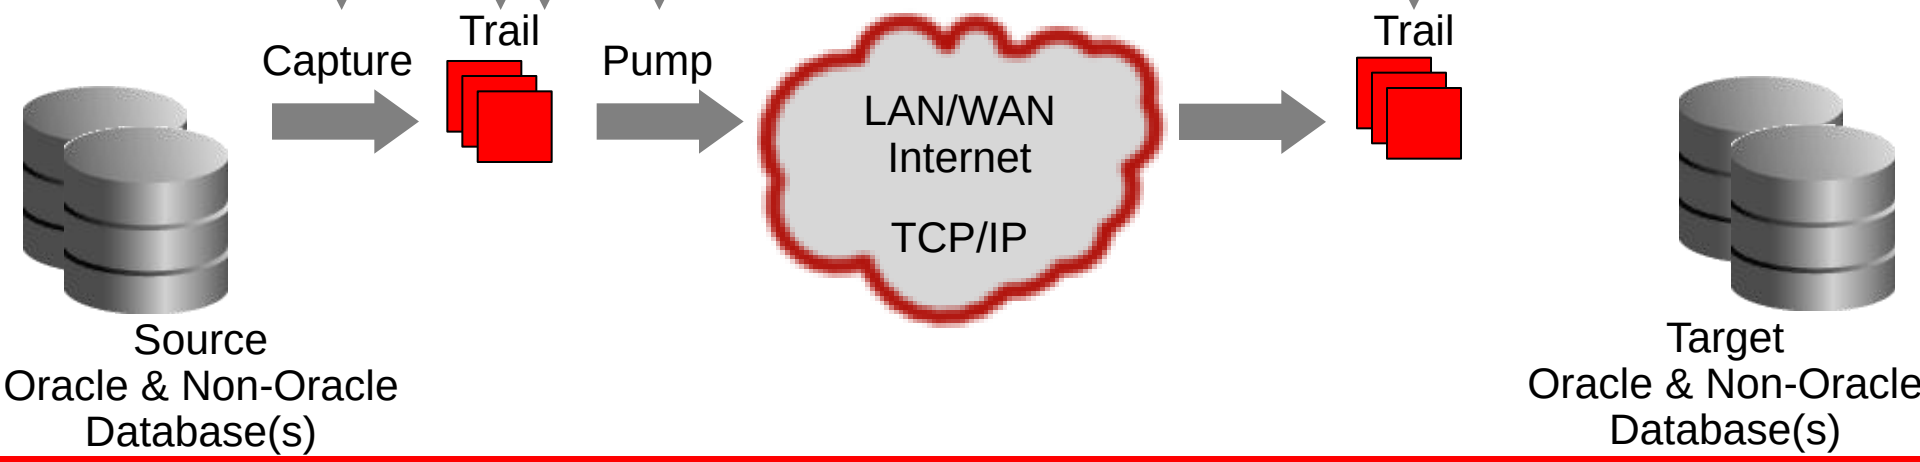

How Oracle GoldenGate Works

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.

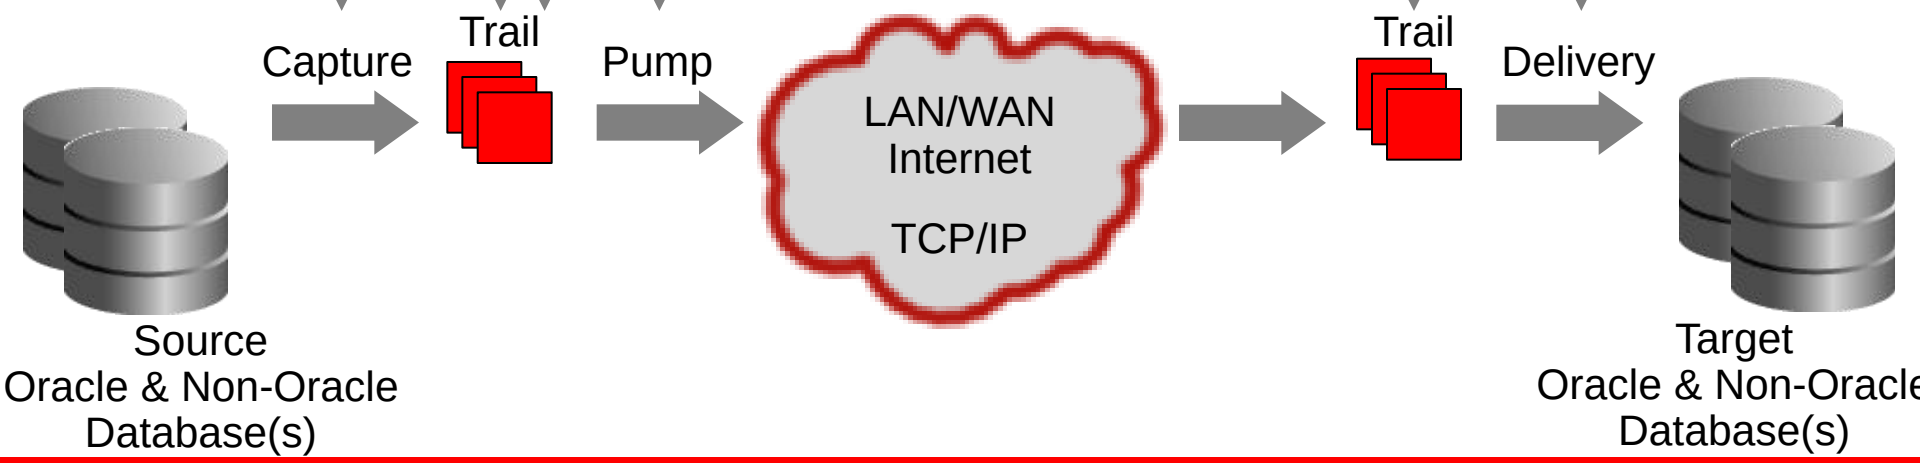

How Oracle GoldenGate Works

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.

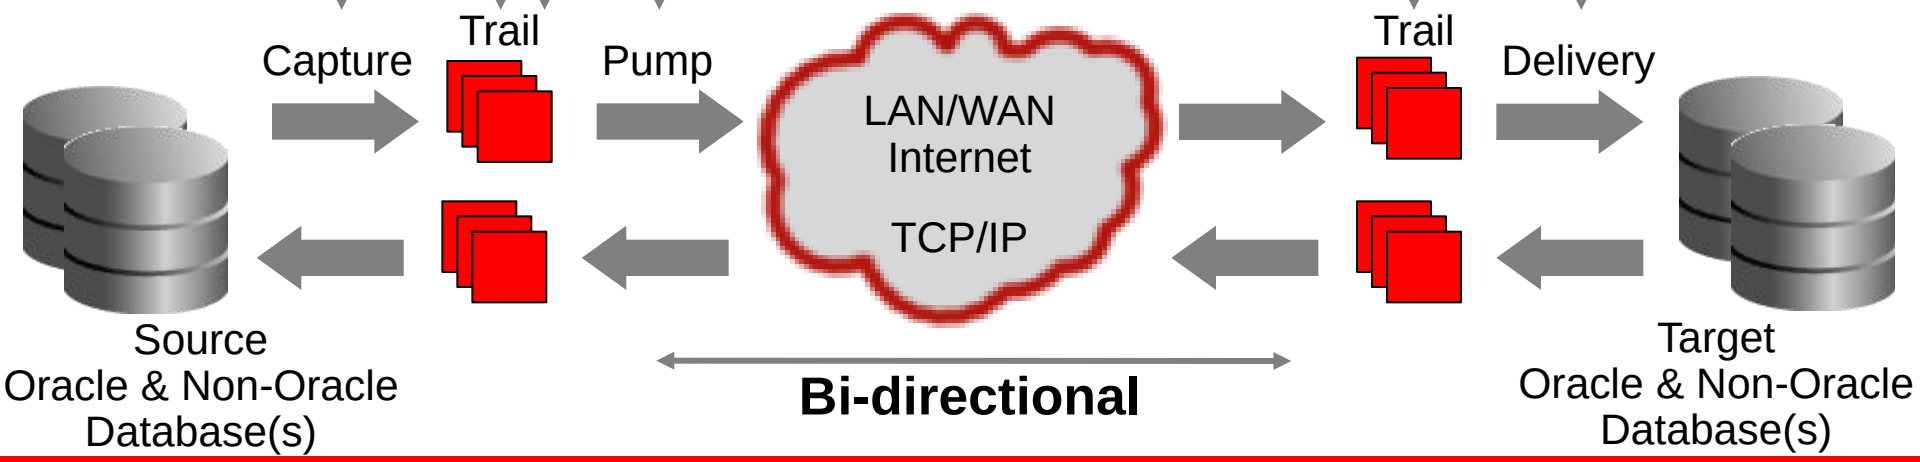

GoldenGate Checkpointing

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

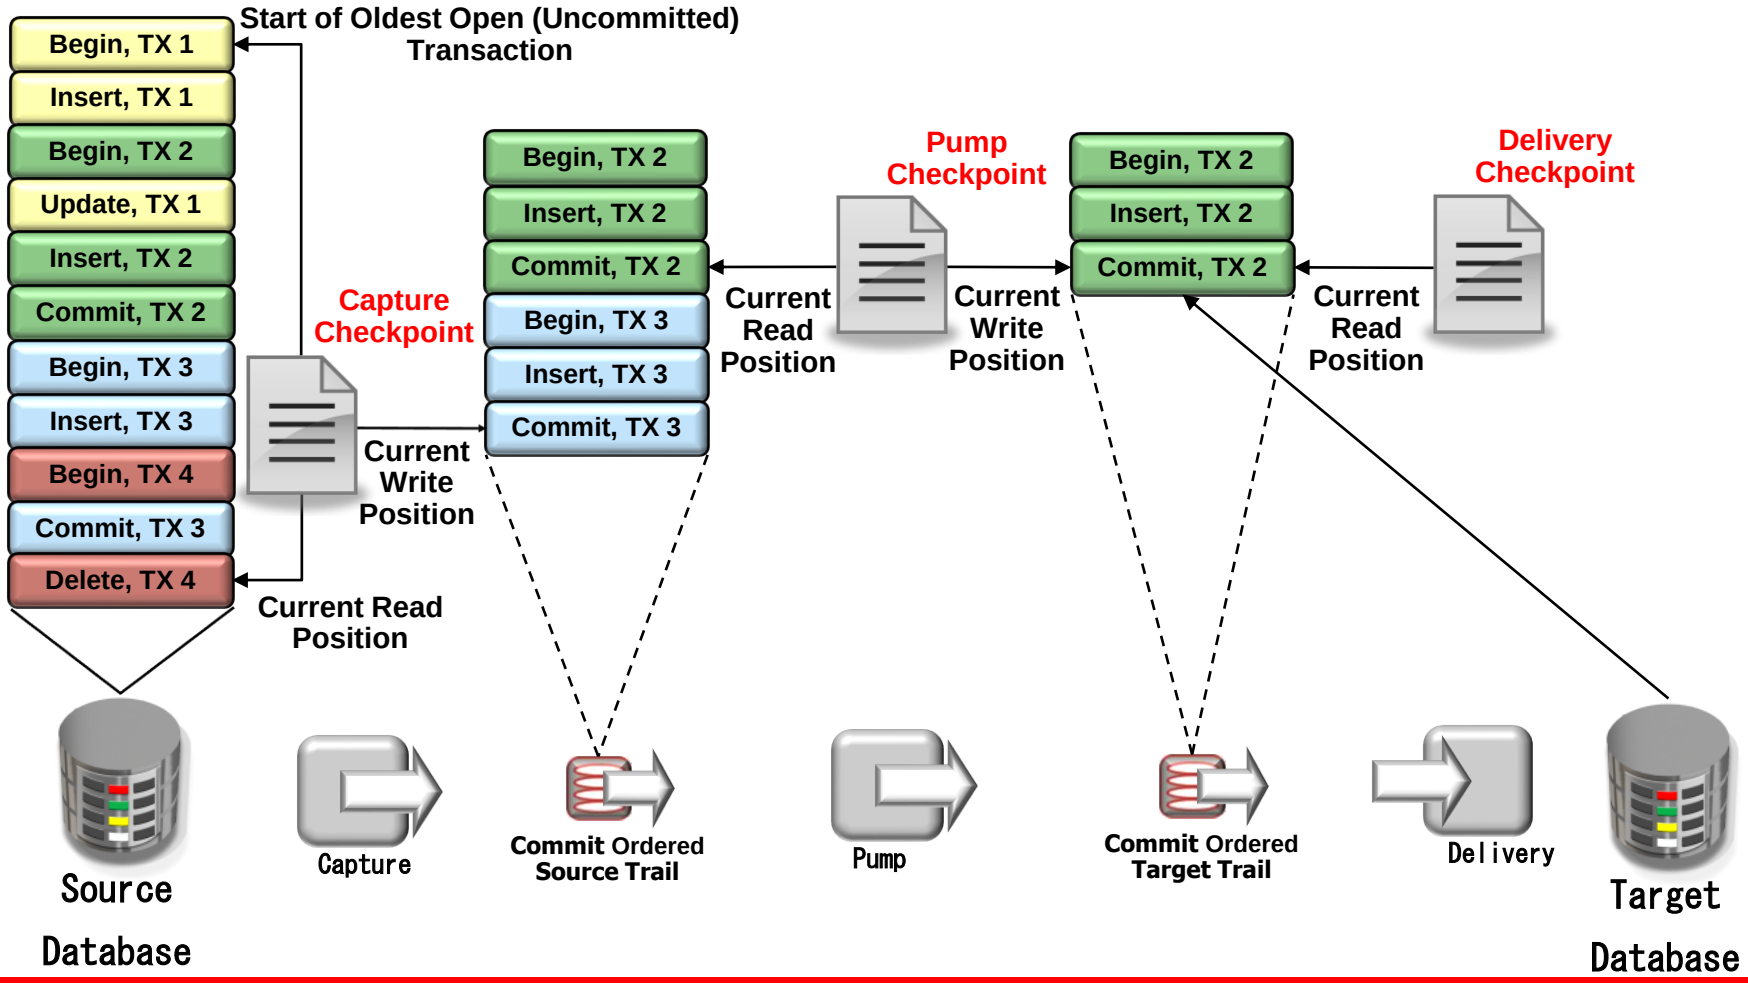

GoldenGate – Scaling for Performance

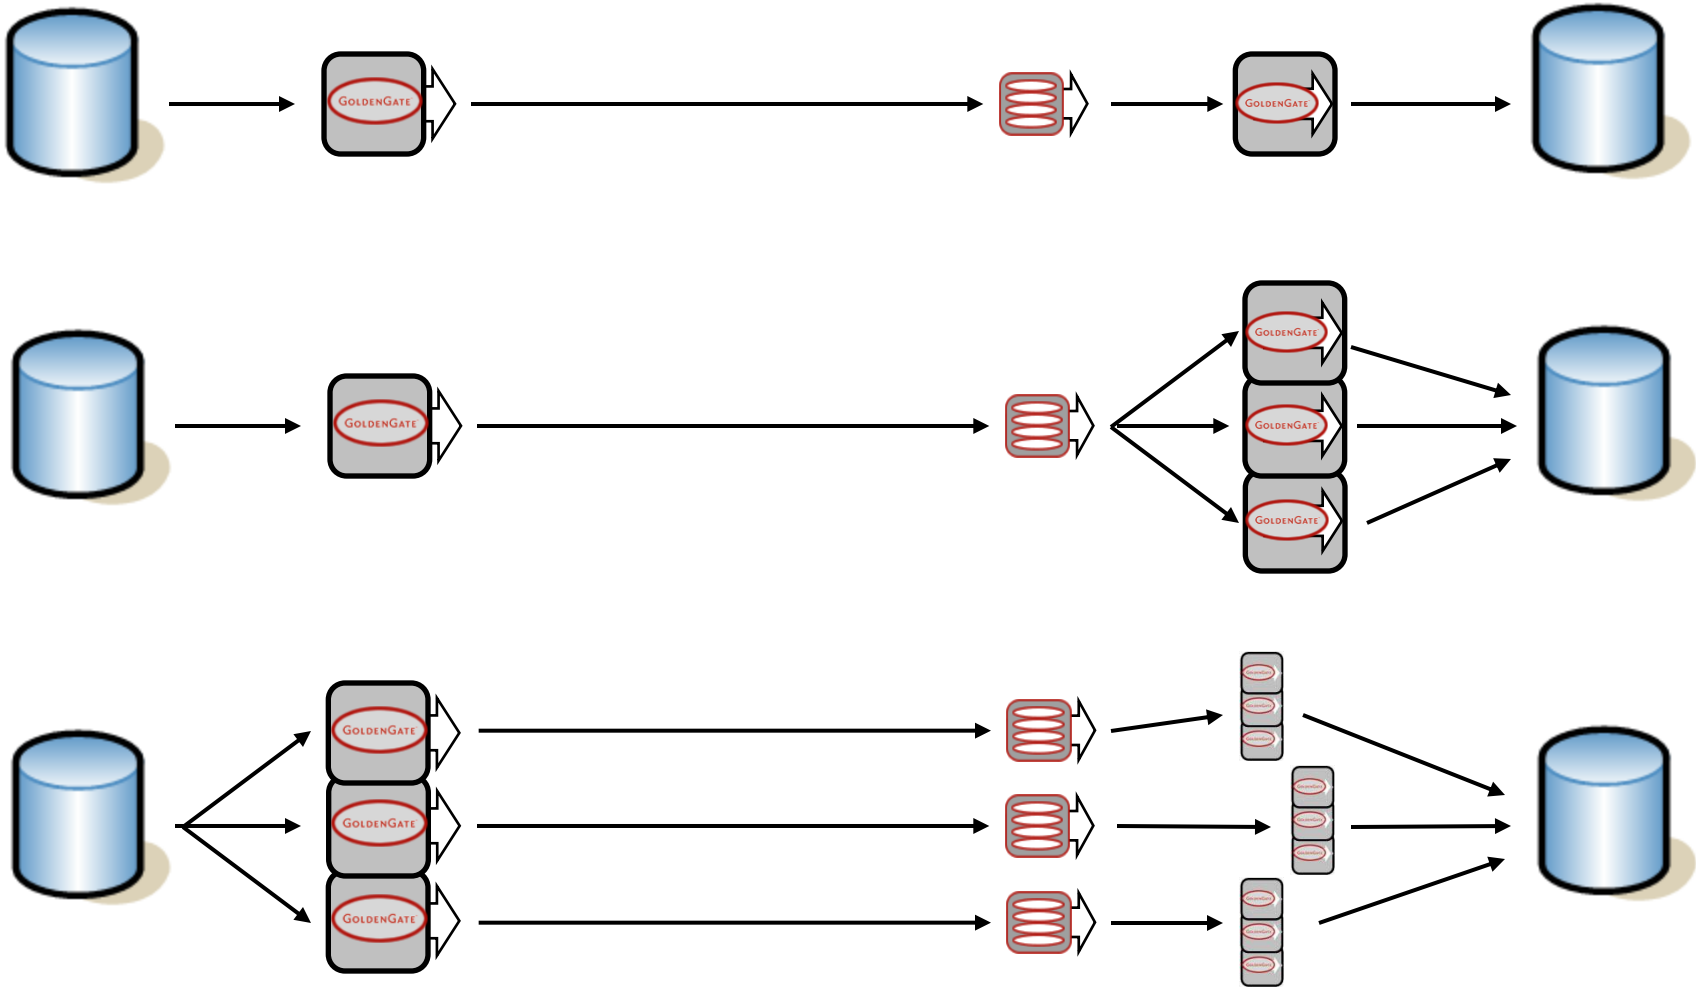

Zero Downtime Oracle Upgrade Implementation Steps:

Example of 9*i* → 11g Cross-Platform

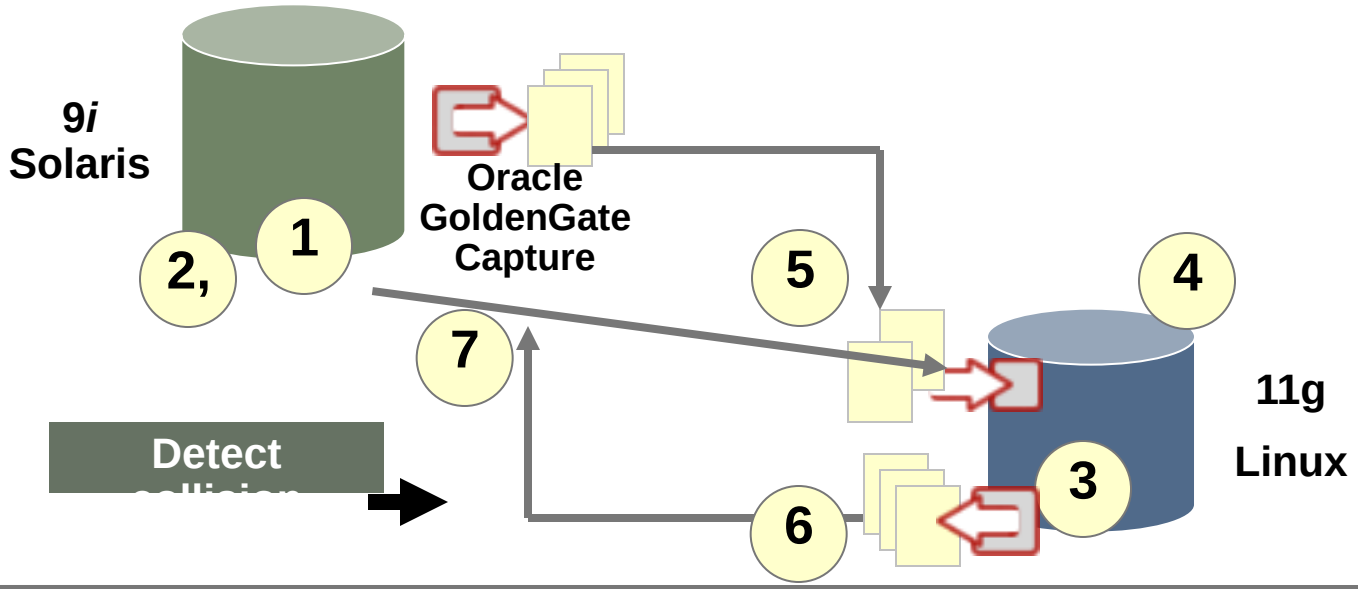

1. Start Oracle GoldenGate Capture module
2. - 4. Initial loading, export import of a new 11g target db (ELT/flat files/jdbc/native db loaders/import export tablespaces etc.)
5. Start Oracle GoldenGate Delivery module at target
6. Start Oracle GoldenGate's Capture at 11g
7. Start Oracle GoldenGate's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

Oracle GoldenGate 11g: Heterogeneity

Databases	O/S and Platforms
Oracle GoldenGate Capture: ▪ Oracle ▪ DB2 NEW for v 9.7 ▪ Microsoft SQL Server ▪ Sybase ASE ▪ Teradata ▪ Enscribe ▪ SQL/MP ▪ SQL/MX ▪ MySQL NEW ▪ JMS message queues NEW **Oracle GoldenGate Delivery:** ▪ All listed above, plus: ▪TimesTen, DB2 for iSeries NEW ▪ Exadata, Netezza, Greenplum, and HP Neoview	Linux Sun Solaris Windows 2000, 2003, XP HP NonStop HP-UX HP OpenVMS IBM AIX IBM z Series zLinux

Customer Example: Zero Downtime Migration

eDialog

Goals

- 24x7x365 provider of advanced e-mail and multichannel marketing solutions to business worldwide helping marketers transform conversations into conversions.
- Ensure absolute business continuity when migrating data to a new data infrastructure

Solution

- Oracle Exadata as the foundation for new data infrastructure that ensures continuous high-performance marketing services and campaign analysis.
- Used GoldenGate for a phased migration with more than 12 terabytes of data from heterogeneous legacy environments

Return on Investment

- Completed the phased migration in six months
- Gained the ability to complete the migration in phases, enabling e-Dialog to test the new environment over time
- Reduced downtime during the massive migration effort
- Improved throughput by 50% and cut report generation time in half

Customer Example: Real-Time DW on Exadata

AVEA

Goals

- Supporting campaigns management with timely customer information
- Reducing batch windows while data increases and improving the performance of ETL and reporting

Solution

- GoldenGate feeds real-time data from CRM, Billing and other key systems to ODS
- ODI extracts from the ODS and loads near real-time data to Exadata DW
- New solution replaced IBM Infosphere Data Stage
- OBI EE is used for real-time reporting

Return on Investment

- Access to timely data for customer segmentation in the Siebel CRM campaign management system
- Batch window for the DW decreased by 50%
- Number of reports generated from the DW has increased by 10 times


ORACLE®


ORACLE®


ORACLE®


ORACLE®


ORACLE®