



An Oracle White Paper  
July 2013

# Using Oracle GoldenGate to Achieve Operational Reporting for Oracle Applications

Introduction .....	2
Right Time for Reporting .....	3
A Common Solution for Operational Reporting .....	3
Operational Reporting Without Performance Impact .....	4
Understanding Oracle GoldenGate.....	5
Architecture .....	5
Key Features of Oracle GoldenGate Core Platform .....	7
Benefits of Oracle GoldenGate in Operational Reporting Solutions	8
Operational Reporting for Oracle Applications .....	9
Operational Reporting for Oracle Applications Using GoldenGate ...	10
Operational Reporting with Oracle E-Business Suite .....	10
Operational Reporting for Oracle PeopleSoft .....	11
Operational Reporting for Oracle JD Edwards .....	12
Operational Reporting Using BI Publisher with Oracle Siebel CRM	13
Operational Reporting with Oracle Business Intelligence Applications	14
Summary .....	14

## Introduction

Accessing real-time information from Oracle Applications for reporting purposes traditionally comes with a price. Many types of reports have the potential to drive up resource utilization and cost, while driving drive down user response times throughout the enterprise. Many organizations either accept or work around these types of situations. Undesired results may include changing the usability of the application, spoiling a user's experience and putting at risk business-critical operations and revenue-generating activities. Oracle Application users understand this, so in order to give users the ability to collate data into usable forms, IT groups must either scale up existing systems or build data transportation processes to extract and load data into reporting systems. Yet, data transportation processes tend to be batch oriented, and as a result these processes create data that is out of time with the rest of the organization.

In this paper, we discuss an Oracle-certified solution that gives organizations access to real-time data from Oracle Applications including Oracle's Siebel CRM, PeopleSoft, and E-Business Suite applications, for reporting purposes, with virtually no interruption to business processes or impact on the end user experience. Through the real-time, log-based change data capture, routing, transformation and delivery capabilities in Oracle GoldenGate, organizations can deploy a cost-effective alternative to the standard batch oriented data transportation systems. This allows faster and lower-impact access to real-time data. Furthermore, Oracle GoldenGate's heterogeneity provides the ability to leverage lower-cost systems and help reduce overall cost of ownership.

## Right Time for Reporting

Many organizations would like to take advantage of reporting on real-time data from business-critical systems but are concerned about the impact to the end user community of those applications. For example, with Oracle E-Business Suite Financials, end of quarter activities for organizations can be very busy and in some cases, very stressful as organizations attempt to close the books. Users executing reports to view open or completed contracts, while the contracts are being updated, have the potential to slow the process down. This inevitably results in organizations creating solutions that physically move user reporting off the primary system to a separate dedicated reporting instance. Traditionally, organizations focus on selecting solutions that extract data in large batches. Yet there is an ‘Achilles heel’ to batch extracts, as the resource spikes created by those extracts can actually drive up resource usage to the point where application response times to end users are adversely affected. Many solve this by running batch extracts during off hours, but the trade off is that the data available in the reporting systems is often not as timely as desired.

## A Common Solution for Operational Reporting

As mentioned above, to reduce reporting’s impact on production systems, companies move the transactional data to a dedicated reporting server. In this section, we will look at a common approach— extract, transform, load (ETL)— for setting up and feeding either dedicated or centralized data warehouses and why this method may not be adequate.

### ETL-Based Bulk Data Movement

To keep reporting twins or centralized data warehouses in sync, organizations often move data from transactional sources using ETL approaches. This involves:

- First, using SQL, columns of the source application database are queried to determine if a row was created or updated during a particular time period.
- Second, as part of an ETL process, the changed data is extracted, exported to a flat file, and loaded into staging tables on the reporting twin or data warehouse. In many cases, the tables in the twin are loaded directly, ending the process.
- Third, a process applies the data changes to the target system’s reporting tables.

Unfortunately, this approach does not provide the real-time data required for true operational reporting, requires specific windows during off-business hours to move data and can still cause significant overhead on source systems, which can negatively affect application usability. Here are the main issues with this approach:

- Batch windows: With ETL integration, the data transfer processes are executed during maintenance windows when the data sources are quiesced; this is necessary to ensure that data sources don’t change during data acquisition, which would create inconsistencies. For

business-critical systems these batch windows are shrinking while the business operations approach 24/7. Further, handling ever-growing data volumes in these shrinking batch windows is becoming a major challenge for ETL users.

- **Overhead on the source:** Source tables are queried and potentially large amounts of data are copied either to other tables for export or directly to flat files in the file system. The net is that the more data extracted, the more resources were required to perform the work.
- **Reliability/recoverability:** Because ETL tools are geared towards processing data in batches, they do not effectively maintain the transaction integrity of data in transit. If an ETL process is interrupted, the partially transferred data often cannot be applied.

To decrease data latency, some ETL products can perform—or be customized for—change data capture (CDC). However, to identify changed data, ETL tools need to store additional data in source tables, such as timestamps. Most databases were not designed to accommodate these extensions, and making changes to the database schema can create issues for the source applications. Such configurations also may place a burden on production systems, because complex queries would need to be run frequently throughout the day across the entire database to identify changes. When the main goal is to reduce impact on production systems, this approach fails to supply timely data without performance degradation.

How can an organization provide real-time, cost-effective, customizable reporting instances, which do not impact business critical systems?

## Operational Reporting Without Performance Impact

There is an alternative to using batch processing or physical replication solutions to replicate large volumes of data. Organizations do not have to limit users' options for reporting, nor do they need to create mirror images of their massive production systems.

### **Low-Impact, Real-Time Data Integration**

Organizations need an unobtrusive solution that can synchronize the desired data between transactional and reporting systems in real time, yet with minimal impact on the performance of the source and target databases. Log-based real-time data integration delivers timely data with virtually no system impact and without requiring intrusive application modifications. In addition to allowing organizations deploy reporting systems side-by-side with the transactional systems, it also allows organizations to capture once and replicate to many other target systems.

By using real-time data integration for offloading operational reporting to a dedicated reporting server, organizations can free up resource usage from business-critical systems and optimize the reporting environment using a wide variety of hardware and operating system combinations. With this method, organizations can utilize one or more replicas, either full or partial, tuned exclusively for reporting.

In this paper we will focus on Oracle's real-time data integration product Oracle GoldenGate and its certified use in operational reporting solutions for Oracle applications.

## Understanding Oracle GoldenGate

Oracle GoldenGate offers a real-time change data capture and replication software platform for transactional data management to meet the needs of today's transaction-driven applications. The software provides guaranteed capture, routing, transformation, and delivery of data across applications and environments in real time. Using this technology, customers can efficiently deploy the latest features and functionality of critical business applications without the risk of standard big-bang upgrades.

### Architecture

Oracle GoldenGate leverages a decoupled architecture to move high volumes of changed data between heterogeneous databases with subsecond latency while preserving transaction integrity. The application can be configured to enable a variety of solutions for continuous availability, disaster tolerance, and real-time data integration. As shown in the diagram below, the Oracle GoldenGate architecture consists of three distinct components Capture, Trail Files, and Delivery. GoldenGate's processes can perform their tasks independently to facilitate rapid, seamless data replication, which is critical to support application-version coexistence.

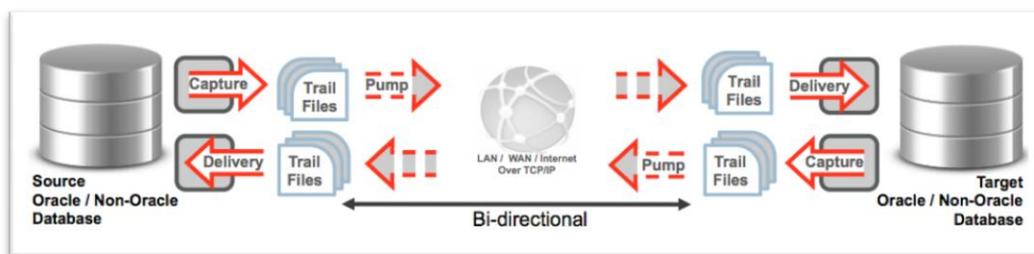


Figure 1. The Oracle GoldenGate architecture supports a variety of topologies, including bidirectional configurations.

### Oracle GoldenGate Capture

The Oracle GoldenGate Capture module resides with the source database and looks for new transactional activity by monitoring database transaction (redo) logs for the results of insert, update, and delete operations. When a change is detected, it is immediately captured for distribution. Capture supports a wide range of database versions including Oracle Database, Microsoft SQL Server, IBM DB2 mainframe and LUW, Sybase ASE, Enscribe, SQL/MP, SQL/MX, and Teradata. Oracle GoldenGate is also certified to support Oracle Exadata and, when used with Oracle Application Adapters 11g, can capture from JMS-based messaging systems.

The Oracle GoldenGate Capture module moves only committed transactions—intermediate activities and rolled-back operations are filtered out—which not only reduces infrastructure load but also eliminates potential data inconsistencies. Further optimization is achieved through transaction grouping and optional compression features. With Oracle GoldenGate 11gR2, Oracle Database users can also take advantage of Integrated Capture, which relies on Oracle’s internal log parsing and processing to capture DML transactions. Integrated Capture allows GoldenGate to support all flavors of compression used by Oracle Database and Oracle Exadata.

### **Oracle GoldenGate Trail Files**

Trail Files, an integral part of Oracle GoldenGate’s proprietary queuing mechanism, store the changed data in a transportable, platform-independent universal data format. Trail Files reside on the source and target server but exist outside of the databases to ensure heterogeneity, improve reliability, and minimize data loss. This architecture reduces the impact on the source system because no additional tables or multiple queries to the database are required to support the capture processes. The Capture module reads once and immediately moves the captured data to the external Trail File for delivery to the target(s).

In the event of an outage at the source or target, the Trail Files contain the most recent changed data up to the point of the outage, and the changes are applied once the systems are back online.

### **Oracle GoldenGate Delivery**

The Oracle GoldenGate Delivery module takes the data transactions from the latest Trail File and applies that data to the target using the native SQL for that relational database management system—delivery can be made to any ODBC compliant database. The Delivery module applies each transaction in the same order as it was committed and within the same transactional context as at the source, to ensure consistency and referential integrity at the target. Delivery uses a number of techniques to optimize the application of data to the target. Changed data can also be provided as a flat file to integrate with third-party ETL products. Oracle GoldenGate can format text in any way, including, but not limited to XML and delimited formats, to be published to enterprise messaging systems.

### **Oracle GoldenGate Manager**

To give users control over Oracle GoldenGate processes, Manager provides a command line interface to perform a variety of administrative, housekeeping, and reporting activities, including:

- Setting parameters to configure and fine-tune Oracle GoldenGate processes
- Starting, stopping, and monitoring capture and delivery modules
- Critical, informational event, and threshold reporting
- Resource management
- Trail File management

Manager executes requests on demand as well as unattended. For example, it can be used to restart Oracle GoldenGate components as well as monitor latency. Manager also automatically recycles Trail File data when no longer needed, providing insurance against inadvertent disk-full conditions and offering an alternative to error-prone manual housekeeping procedures.

For enhanced management of Oracle GoldenGate 11g processes and solutions, customers should consider adding the Oracle Management Pack for Oracle GoldenGate. Management Pack for Oracle GoldenGate is a centralized, server-based graphical enterprise application that offers an intuitive way to define, configure, manage, monitor and report Oracle GoldenGate processes.

### Key Features of Oracle GoldenGate Core Platform

Oracle GoldenGate utilizes the following features to facilitate operational reporting without performance impact on source applications:

- **Transformations and mappings.** Oracle GoldenGate can flexibly accommodate transformations and mappings within either the Capture or Delivery modules—no middle tier server is needed. The product supports table and row filtering based on user-defined criteria. Explicit mapping and transformation rules can be applied via built-in functions, user-supplied code, and stored procedures. Those rules may range from simple column assignments to more complex transformations, for which Oracle GoldenGate provides a suite of date, math, string, and utility functions.
- **Flexible topology support.** Oracle GoldenGate’s architecture allows customers to support a variety of topologies, including one source to one target, one-to-many, many-to-one, many-to-many, and cascading, and bidirectional configurations. For example, Oracle GoldenGate allows a configuration in which a second Capture component, called a “Data Pump,” continuously pushes the Trail Files from the source system to multiple target systems.
- **Bidirectional configuration support.** Oracle GoldenGate enables both active-passive and active-active solutions for maximum system availability.
- **Conflict detection and resolution.** Bidirectional, active-active implementations require conflict detection and resolution capabilities because multiple systems are actively processing and sharing database transactions. Oracle GoldenGate provides conflict detection and resolution options that can be implemented globally, object-by-object, based on data values and complex filters, or through event-driven criteria.
- **Routing and Compression.** Oracle GoldenGate utilizes TCP/IP for sending data so no geographical distance constraints are imposed between source and target systems. In addition, Oracle GoldenGate can apply additional compression to the data as it is routed.
- **Data Encryption.** Data encryption ensures secure, confidential data transmissions.

## Benefits of Oracle GoldenGate in Operational Reporting Solutions

Oracle GoldenGate offers the following benefits that are not possible with running operational reporting on the production system:

- **Workload and Resource Utilization.** Rather than all users generating reports and running resource intensive operations against the single primary system, organizations can split operations to systems minimize the overhead on the production environment.
- **Optimization.** Reporting systems operate more efficiently if they are tuned specifically for the usage. Large sort areas, partitions, and different types of indexing are better suited for reporting versus transactional systems.
- **Full Use of the Reporting Instance.** Oracle GoldenGate does not limit the type of operations that can be executed on the reporting instance. Reporting needs that include the creation of temporary tables or other processes to prepare data for reporting can be run on the reporting instance without impacting primary business systems.
- **Transformations using Oracle Data Integrator.** For use cases with complex data transformation requirement, Oracle GoldenGate can be combined with Oracle Data Integrator. Oracle Data Integrator (ODI) offers bulk data movement and transformations with an extract, load, and transform (E-LT) architecture and offers tight integration with Oracle GoldenGate. In this integrated solution, GoldenGate moves change data to the staging area of the target system in real-time and ODI transforms and loads data in micro-batches within the target database for user tables. This solution leverages GoldenGate's log-based capture to avoid impact on the source systems, and minimizes batch windows with micro-batch data processing.
- **Expansion Options.** Operational Reporting solutions are by design read-only and tend to be deployed unidirectionally (i.e. data flows from the source transactional system to the reporting instance, where reports are executed), but by using Oracle GoldenGate the option exists to include bi-directional replication for key system tables or for entire sets of data. Oracle GoldenGate is flexible to fit the needs of both the application being used and business users.
- **Reporting Solution Options.** Operational Reporting using Oracle GoldenGate software platform for data integration provides organizations with the option to use either reporting tools such as Oracle's BI-Publisher or tools provided in Oracle's applications.
- **Certified by Oracle.** Oracle GoldenGate is tested and certified on major Oracle applications including E-Business Suite, JD Edwards, PeopleSoft and Siebel CRM. Development teams for these applications validated that Oracle GoldenGate does not adversely impact or risk the deployments of those Oracle applications.

## Operational Reporting for Oracle Applications

As highlighted above, Oracle GoldenGate's flexible architecture is certified to support operational reporting, by replicating business data to a secondary system which would be used to execute read-intensive operations, such as reporting. (See figures 3.0 and 3.01 below for a typical operational reporting solution and configuration.)

Oracle GoldenGate reads changed data from database transaction logs rather than from the database tables themselves. Because it requires minimal modifications to the applications, this solution provides organizations with a compact, non-intrusive, and easily configured method for providing access to real-time data for reporting purposes. As shown below, organizations can replicate data to create highly optimized reporting solutions.

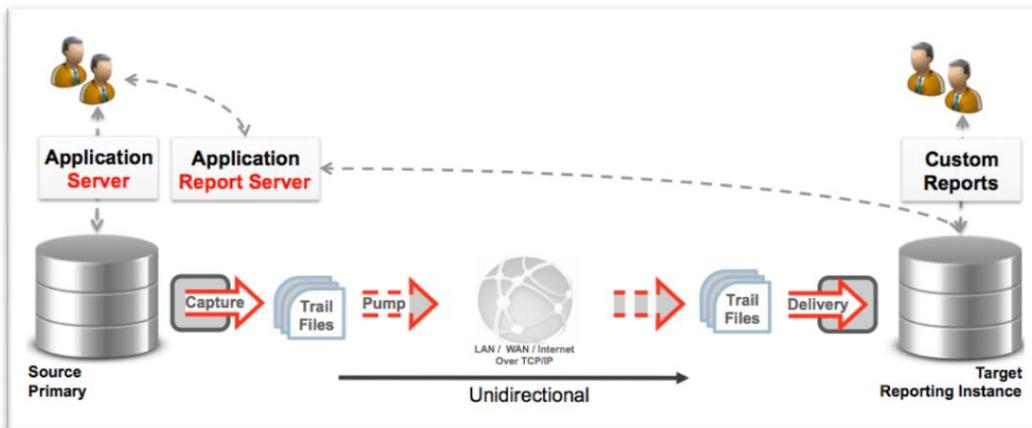


Figure 2.0 Oracle GoldenGate Operational Reporting Solution

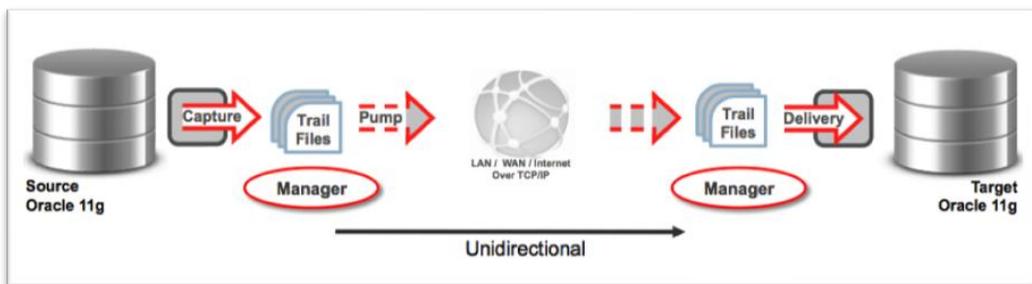


Figure 2.1 Oracle GoldenGate Operational Reporting Conceptual Configuration (Unidirectional)

Taking advantage its flexibility, Oracle GoldenGate can be configured to support the needs of Oracle Applications to provide the means to move data bi-directionally. That allows for seamless integration of the built-in reporting tools to use the reporting instance.

## Operational Reporting for Oracle Applications Using Oracle GoldenGate

Oracle GoldenGate complements Oracle applications to enable execution of reports from within the application, such as E-Business Suite, or using Oracle's reporting tools, such as BI-Publisher.

Below are examples of Oracle GoldenGate's certified use case descriptions for the major Oracle applications:

### Operational Reporting with Oracle E-Business Suite

The flexible architecture of Oracle GoldenGate can be adapted to support operational reporting for Oracle E-Business Suite. The solution replicates transactional data from the E-Business Suite database to a reporting instance with sub-second latency, which would be used to execute read intensive real-time reporting. Additional details can be found in Knowledge Document 1112325.1 on Oracle Support.

The figure below is a conceptual illustration of a typical operational reporting solution for Oracle E-Business Suite.

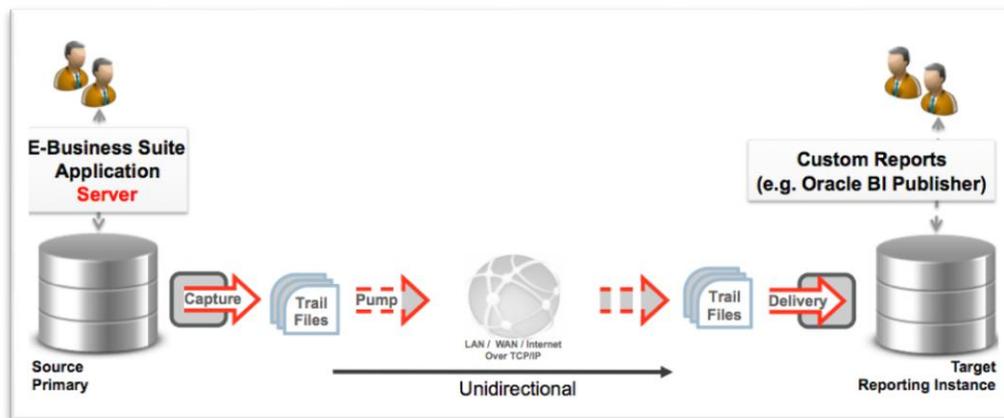


Figure 3.0 – Oracle E-Business Suite Operational Reporting

Oracle's own implementation of Oracle GoldenGate for Oracle E-business Suite illustrates the dramatic improvements in data freshness and report performance that can be achieved. In 2011, the decade old contracts reporting solution running within Oracle's global single instance of E-Business Suite had reached a breaking point. Service contracts represent over \$20 billion of revenue per year for Oracle, with well over 100,000 renewals and new contracts to be booked every quarter. The service contracts business generates up to 800,000 changed rows of data in Oracle's global single instance (GSI) of E-Business Suite every day, with daily volume spikes as quarter end approaches. As this system continued to grow, it became increasingly less useful, as

stale data began to deprive Oracle of a good day-to-day picture of a key component of its business.

As Oracle's installed base increased, twice daily refreshes were no longer possible. Instead, sometimes data was more than 48 hours old. On top of this, reports were routinely taking longer than eight hours to complete, and especially long reports sometimes had to be terminated after 12 hours in order to keep the entire system from simply grinding to a halt. Operations and tactical decisions were being hampered by outdated information. A drastic change was required.

Like most customers, Oracle considered a variety of options before settling on a solution that would allow only the relevant subset of changed data to be replicated in near real-time using Oracle GoldenGate. With this solution, the Operation Data Store was designed to contain only the subset of source system data required for operational reporting.

By using Oracle GoldenGate on E-Business Suite to drive Oracle's service contracts operational reporting solution, data freshness has improved by orders of magnitude. Data freshness is now measured in seconds, and even very large reports now run in seconds or minutes rather than hours. These improvements in have already driven a significant increase in usage, and the ODS solution is now being expanded from service contracts into other areas of Oracle's business.

### Operational Reporting for Oracle PeopleSoft

Oracle GoldenGate's flexible architecture is certified to support PeopleSoft operational reporting, to replicate business data to a secondary system which would be used to execute read-intensive operations, such as reporting. The solution is designed for all of the tables within application schemas to be replicated for the purpose of off-loading reporting functions to a secondary / non-OLTP reporting instance. Because Oracle GoldenGate requires minimal modifications to the applications, this solution provides organizations with a compact, non-intrusive, and easily configured method for providing access to real-time PeopleSoft data for reporting purposes.

Oracle GoldenGate supports two different options for report offloading for Oracle PeopleSoft. With option one, the replica can be either Oracle or non-Oracle, schemas can be different, and third party tools can be used for reporting.

With option two, or integrated operational reporting, reports can be created to run via the PeopleSoft Process Scheduler, which is pre-configured to point to the reporting instance.

The figure on the next page is a conceptual illustration of a typical integrated operational reporting solution for Oracle PeopleSoft.

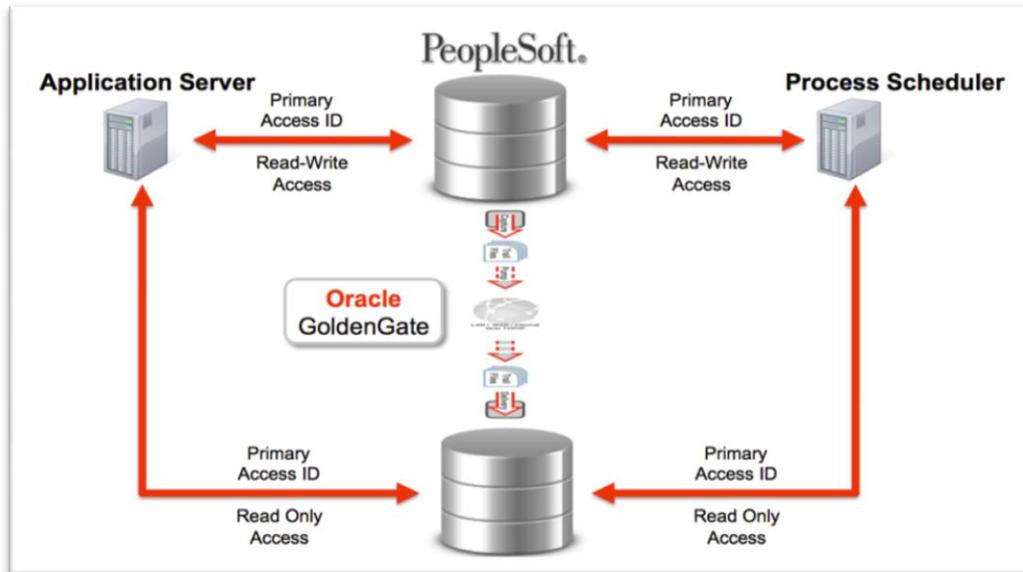


Figure 4.0 – Integrated Reporting for Oracle PeopleSoft using Oracle GoldenGate

This second integrated reporting option is an out of the box solution that was developed by PeopleSoft for PeopleTools 8.52 using the Oracle Database. Reports that are executed via the PeopleSoft Process Scheduler should be read-only, but can create temporary tables if needed. In order for Process Scheduler to work seamlessly, configuration of Oracle GoldenGate to include key tables for bi-directional replication is required along with configuration changes to the set-up of the Process Scheduler. Additional details on both operational reporting and integrated operational reporting can be found in Knowledge Document 1114746.1 on Oracle Support.

### Operational Reporting for Oracle JD Edwards

Operational reporting for JD Edwards replicates all tables from the primary JDE database unidirectionally to a reporting instance. Using a second JD Edwards application server, users can access the reporting instance to generate resource intensive reports. The key limitation with this type of configuration is that users should not create data in the reporting instance, as doing so would create an out-of-sync situation between the source and target databases, which would require the reporting instance to be re-initialized.

As JD Edwards has the ability to run on other relational databases, such as Microsoft's SQL Server and IBM DB2 on System i, organizations can take advantage of the heterogeneous capabilities of Oracle GoldenGate to deploy the operational reporting solutions using Oracle Database as the target. Additional details can be found in Knowledge Document 1112406.1 on Oracle Support.

Below diagram shows how Oracle GoldenGate can compliment JD Edwards.

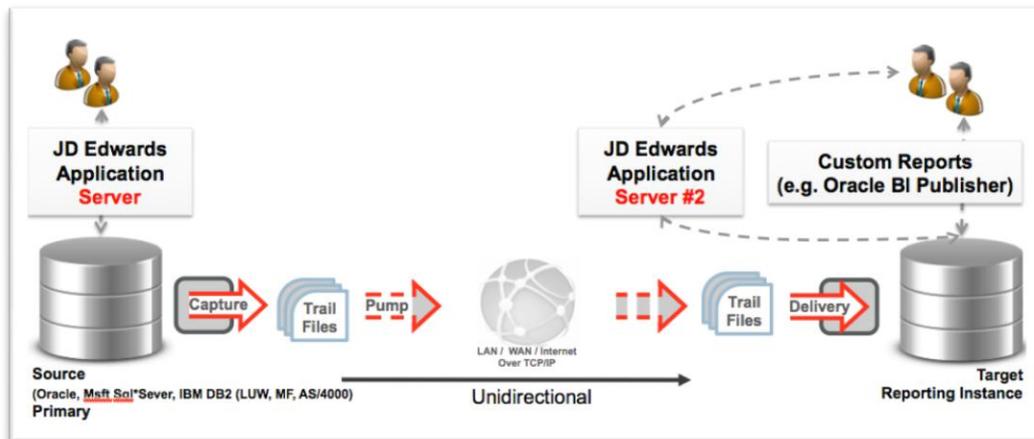


Figure 5.0 – Oracle GoldenGate for JD Edwards Configuration

### Operational Reporting Using BI Publisher on an Oracle Siebel CRM Reporting Instance

Deploying Oracle GoldenGate for operational reporting allows customers to combine technologies to meet the needs of the business. Operational reporting using Oracle GoldenGate is flexible enough to allow reporting tools, such as Oracle's BI Publisher, to be run against the reporting instance. There are very few limitations on what tools can be used to query data from the reporting instance, yet typically the biggest hurdle for business users is to make sense of the complex application schema well enough to construct custom reports.

Oracle GoldenGate provides the means to remap data from the source system to different but like schemas that may be more easily understood by business users. For example, Oracle Siebel CRM stores account details in a group of entities but primarily in a table called S\_ORG\_EXT. Using the flexible configuration options of Oracle GoldenGate, data from S\_ORG\_EXT could be remapped to a table in the reporting instance called ACCOUNTS. Add in remapping of columns from S\_ORG\_EXT to the new ACCOUNTS tables and business users could more easily navigate the schema for creating custom reports. Additional details can be found in Knowledge Document 1112403.1 on Oracle Support.

Below diagram show Oracle BI Publisher for reporting against real-time data from Siebel CRM

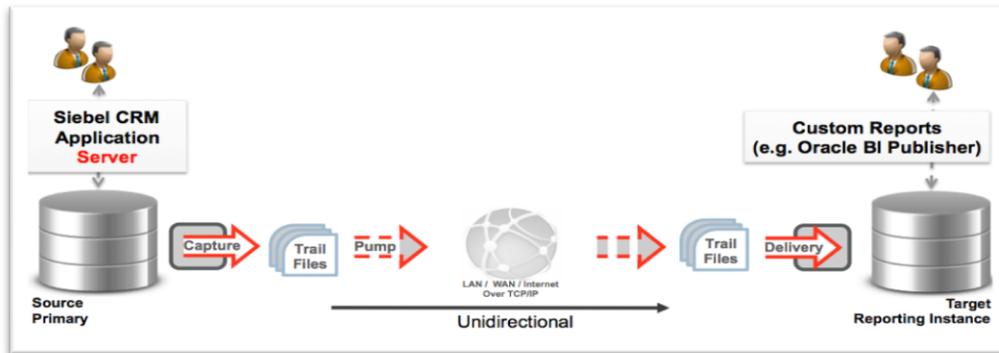


Figure 6.0 – Oracle GoldenGate with BI Publisher Environment Summary

## Operational Reporting with Oracle Business Intelligence Applications

Oracle Business Intelligence Applications 11g, Release 11.1.1.7.1, has been designed to leverage the advantages of Oracle Data Integrator Enterprise Edition (ODI). This new release of BI Applications enables customers to increase IT efficiency and reduce costs with a comprehensive data integration platform that covers all data integration requirements – including big data, application integration, as well as BI / data warehousing.

ODI helps to integrate data end-to-end across the full BI Applications architecture, supporting capabilities such as data-lineage which helps business users identify report-to-source capabilities. In addition, customers can choose the option to replicate their data in real-time using Oracle GoldenGate. Oracle GoldenGate enables real-time business intelligence for improved business insight, query offloading to maximize OLTP performance, zero-downtime data migration, disaster recovery, and active-active database synchronization for continuous availability.

## Summary

Organizations need a low-impact, cost-effective solution to replicate data out of business critical Oracle Applications without interrupting users. Utilizing the flexible architecture of Oracle GoldenGate, organizations can deploy certified solutions for real-time, low-impact, and non-intrusive replication of data from Oracle Applications to systems designed exclusively for reporting. It can also send that data to other systems within the organization that require real-time access to application data locally.

Using the certified operational reporting solution with Oracle GoldenGate, companies can use up-to-the-second transactional data from their Oracle applications without impacting business operations or end user experience.



Using Oracle GoldenGate to Achieve  
Operational Reporting for Oracle Applications  
July 2013

Oracle Corporation  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

[oracle.com](http://oracle.com)



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0113

**Hardware and Software, Engineered to Work Together**