



An Oracle White Paper
Updated September 2013

Using Oracle GoldenGate 12c for Oracle Database

Executive Overview	2
Introduction	3
Architecture Overview	4
Oracle GoldenGate Capture	5
Oracle GoldenGate Trail Files	6
Oracle GoldenGate Delivery	7
Oracle GoldenGate Manager	9
Associated Products	9
One Platform, Many Solutions	11
Zero Downtime Migrations and Upgrades	12
Query Offloading	12
Disaster Recovery and Data Protection	13
Active-Active Database Replication	13
Operational Reporting and Real-Time Data Warehousing	13
Data Distribution and Synchronization for OLTP Systems	14
Oracle GoldenGate for Oracle Database	15
Capture (Extract)	15
Routing (Data Pump)	16
Delivery (Replicat)	17
Transformation	19
Customer Examples	22
Key Benefits of Using GoldenGate for Oracle Database	24
Summary	25

Executive Overview

Business applications have become increasingly critical for transaction processing. As a result, end users must access, analyze, act on, integrate, store, and verify transactional data faster than ever—often in real time—and without system interruption or downtime. This demand for low-latency data is compounded by the exponential growth in transactional data volumes and an increasingly heterogeneous enterprise IT environment, creating a need for data integration and replication solutions that are easy to implement and have little to no impact on business-critical applications.

Oracle GoldenGate is used by major Fortune 500 companies and other industry leaders worldwide for real-time database synchronization to support mission-critical systems to achieve maximum availability and real-time integration in heterogeneous environments. Written for business project owners, key stakeholders, and the entire IT organization, this white paper provides a detailed look at Oracle GoldenGate 12c, its underlying technology architecture and how it works with the Oracle Database.

Introduction

Oracle GoldenGate 12c offers a real-time, log-based change data capture (CDC) ¹ and replication software platform to meet the needs of today's transaction-driven applications. The software provides capture, routing, transformation, and delivery of transactional data across heterogeneous databases in real time. Using this technology, customers can achieve continuous availability for critical systems and real-time data integration for fast, easy access to current and accurate data. See Table 1.

TABLE 1. ORACLE GOLDENGATE KEY FEATURES AND DIFFERENTIATORS

FEATURE	DETAIL
Real-time data feeds	Provides continuous capture and delivery of data from sources to targets with end-to-end subsecond latency. Operates with high performance and low overhead even at high volumes.
Heterogeneity	Captures and delivers data between a variety of relational, open systems/open source, and legacy databases on all major platforms. Captures from, and delivers to, Java Messaging Service (JMS) based messaging systems.
Transactional integrity	Maintains the reliability and accuracy of transactional data as it is moved between systems by enforcing ACID properties and referential integrity.

¹ Oracle GoldenGate's Change Data Capture capability should *not* be confused with the Oracle Database feature Oracle Change Data Capture (Oracle CDC).

Architecture Overview

The Oracle GoldenGate 12c architecture consists of decoupled modules that can be combined across the enterprise to provide maximum flexibility, availability, and performance. This architecture facilitates the movement of transactional data in four simple, yet powerful steps.




Figure 1. Oracle GoldenGate Architecture

Key Operational Steps

- **Capture.** Oracle GoldenGate captures changed data operations committed in the database transaction logs in a nonintrusive, high-performance, low-overhead implementation.
- **Route.** Oracle GoldenGate can use a variety of transport protocols, and it can compress and encrypt changed data prior to routing.
- **Transform.** At any point prior to applying the data to the target system, Oracle GoldenGate can be used to execute a number of built-in functions, such as filtering and transformations.
- **Deliver.** Oracle GoldenGate applies the changed transactional data to one or more database targets with subsecond latency, preserving transactional integrity. Transactional data can be delivered via Open Database Connectivity–compliant databases or through a specialized adapter to a JMS message queue or topic.

Key Components

Oracle GoldenGate consists of four distinct modules and components:

- Oracle GoldenGate Capture
- Oracle GoldenGate Trail Files
- Oracle GoldenGate Delivery
- Oracle GoldenGate Manager

Oracle GoldenGate Capture

The Capture module processes committed transactions resulting from insert, update, and delete operations executed against a database, as well as selected DDL operations, and routes them for distribution. Capture process executable name “Extract” is typically used in documentation to refer to the Capture module of the product.

High-Speed, Low-Impact Data Capture

The Capture module does not require any changes to be made to the source database or the application it supports. To maintain optimal performance, the Capture module employs a range of change capture techniques against the source database. For instance, in databases that include transaction logs, such as Oracle Database (including Oracle Exadata); Microsoft SQL Server; IBM DB2 (LUW), System z and System i; Sybase ASE; MySQL; and those running on HP NonStop/Enscribe, SQL/MP, and SQL/MX, changes are captured through direct access to native database transaction logs (redo logs, if applicable). For Teradata at the source, custom APIs have been developed to allow Oracle GoldenGate to capture committed transactions with the same efficiencies.

Transaction logs contain all changes made to the database and are automatically maintained by the database application independently of Oracle GoldenGate. Consequently, no additional tables are required to run the Capture module, and overhead is greatly reduced as compared with trigger-based capture techniques. The Capture module can automatically adjust its transaction memory based on the size and number of the transactions it is capturing, which optimizes memory usage, allowing even lower overhead on the source systems. When used with Oracle GoldenGate Application Adapters 11g, the product also offers capabilities to capture from JMS-based messaging systems.

Table, Row, and Column Selectivity

When only a subset of changed data from the source needs to be replicated to the target system—such as for real-time reporting purposes—the Capture module allows users to filter tables and rows based on user-defined criteria and ignores the entries in the transaction log that don’t meet the end-user’s needs. Users can optionally select and apply transformation rules to specific columns via built-in Oracle GoldenGate functions, user-supplied code, stored procedures, or Oracle Data Integrator Enterprise Edition.

Efficient Network Use and Large Data Volumes

The Capture module can route transactions over WANs and LANs as well as the Internet, and it can reduce network bandwidth requirements in a number of ways. Typically, the amount of data transmitted is only a fraction of the data that is generated by the database and stored in transaction logs. Because only committed transactions are propagated, intermediate activities and rolled-back operations are not transferred. The Capture module optimizes data transmissions, by bundling individual records into larger, more-efficient packets to avoid potential bottlenecks.

Several levels of data compression are available to further reduce the amount of network bandwidth required for transmission. Depending on data types, data compression can potentially reduce byte transfer by up to 75 percent or more.

For scenarios requiring very large changed data volumes, users can deploy multiple Capture modules to minimize the lag between source and target systems. Additionally, customers running Capture on an Oracle Database can choose to take advantage of Integrated Capture, a multi-threaded capture mechanism that improves performance by interacting directly with a database log mining server to receive data changes. Integrated Capture also provides users with the ability to reduce overhead on the source system by offloading the Capture process to an alternate location, such as the target.

Checkpoints for Reliable Data Delivery

Oracle GoldenGate creates a checkpoint at the last changed transaction whenever a commit boundary is encountered. This enables the delivery of all committed records to the target, even in the event of a restart or cluster failover. Checkpoints store the current position as processed by both the Capture and Delivery modules. Following a network or system outage, Oracle GoldenGate restarts from the last good checkpoint. Oracle GoldenGate also persists uncommitted operations to disk to enable fast and simple data recovery for long running transactions in the event that the replication process is paused or interrupted.

Oracle GoldenGate Trail Files

Trail Files—Oracle GoldenGate’s unique queuing mechanism—contain the most recent changed data in a transportable, platform-independent format called the Oracle GoldenGate Universal Data Format, and can be converted to XML and other popular formats for consumption by different applications. The Capture module can create unique as well as overlapping sets of data in each Trail File. Based on the requirements of the implementation, users can store Trail Files on the target system, the source system, or both. Trail Files can be delivered to alternative queue types and application interfaces.

Flexible, Decoupled Architecture

By staging data in Trail Files, Oracle GoldenGate decouples the data source and target for heterogeneous support. Unlike architectures that implement a tight process-to-process coupling, this decoupled architecture allows each module to perform its tasks independently. In the event of an unplanned outage, decoupling ensures that the unaffected systems continue to operate.

Oracle GoldenGate also provides flexibility in the choice of hardware, operating system, and databases for sources and targets. For maximum flexibility and ease of use, customers can use different versions of Capture, Delivery, and Trail Files in the same implementation.

Data Pumps

Depending on the configuration and environment, it is generally preferable to create Trail Files on the source system and use a separate Oracle GoldenGate feature, called a Data Pump², to continuously push—or “pump”—the Trail Files from the source system to the target system(s). This configuration enhances the fault tolerance and reliability of the overall Oracle GoldenGate environment. In the event of a network failure (between the source and the target systems), Oracle GoldenGate can continue to capture transactions because the data can be queued up locally in the Trail Files on the source, enhancing the recoverability in case of database failures.

The Data Pump feature is strongly recommended if data needs to be distributed to multiple targets from the same source (one to many). Whereas the Capture module can focus solely on capturing transactions, individual Data Pumps can be set up to distribute the data to those targets, increasing the efficiency of the overall environment. Fault tolerance is also greatly increased in such a configuration because any failure associated with one target has no impact on the source capture or delivery to other targets—transactions will continue to be captured, routed, and delivered to the other targets even when one of them is down. Data Pumps can also be used to route data through an intermediate system, even if that system doesn’t have a database installed.

Archival and Audit Capabilities

Trail Files can create an archive of purged information from the source database by transforming delete and update records into inserts in a different location. For auditing and compliance purposes, Oracle GoldenGate can also maintain a separate history table to track each update to individual records as they change.

Oracle GoldenGate Delivery

The Delivery module takes changed transactional data that has been placed in a Trail File and immediately applies it to the target database. Supported target databases include Oracle Database (including Oracle Exadata); Microsoft SQL Server; IBM DB2 (LUW), System z, System i; Sybase ASE; Teradata; those running on HP NonStop/Enscribe, SQL/MP, and SQL/MX, IBM Netezza, Greenplum, MySQL, PostgreSQL, and TimesTen. Through the use of Oracle GoldenGate Application Adapters, Oracle GoldenGate also has the capability to publish changed data to a messaging system in XML or other formats, as well as provide data in flat files

² Oracle GoldenGate’s Data Pump component is *not* related to the Oracle Database utility Oracle Data Pump.

for third-party products, such as an ETL system. Delivery process executable name “Replicat” is typically used in documentation to refer to the Delivery module of the product.

Data Integrity and Transaction Consistency

Each Delivery module applies each database change in the same order as it was committed in the source database to provide data and referential integrity. In addition, it applies changes within the same transaction context as they were on the source system for consistency on the target.

Column Mapping and Transformation

As with Capture, users can configure the Delivery module via user-defined criteria too not only specify target tables but also individual rows and columns. By default, the Delivery module populates any target table column with data from a source table column if the two columns share the same name, and this is also true of like-named tables. However, you can easily configure Oracle GoldenGate to move data from a single table into multiple target tables or vice versa. This can be used to normalize or denormalize data in a data warehouse or OLTP environment.

Users can also define explicit mapping and transformation rules, ranging from simple column assignments to more-complex transformations for which Oracle GoldenGate provides a suite of date, math, string, and utility functions. The module also supports the use of stored database procedures and functions and enables implicit mapping and explicit rules to be combined. If additional transformations, data quality, aggregation, and other functionality are required, Oracle GoldenGate 12c integrates with Oracle Data Integrator Enterprise Edition 12c to support end-to-end data integration.

Optimized High-Speed, High-Volume Data Delivery

The Delivery module provides a variety of techniques to optimize the posting of changed transactions to the target database. Oracle GoldenGate’s posting processes, where possible, run local to the target database, maximizing throughput by avoiding network limitations. The Delivery module also minimizes disk I/O while preserving original transaction properties. In addition, where possible, updates are executed via native database interfaces rather than through middleware, and internal caches are used to ensure fast execution of repetitive statements.

When applying transactions to a target Oracle Database 11.2.0.4 or higher, users can take advantage of Integrated Delivery. Integrated Delivery leverages Oracle Database parallel apply servers for automatic dependency tracking and parallel aware apply. When applying transactions to non-Oracle datastores, Coordinated Delivery capability was introduced with Oracle GoldenGate 12c to simplify deployment of multiple Delivery modules, which can minimize lag time in the event of high data volumes during peak processing times or seasonality. These capture-route-transform-apply processes run continuously, so that the most recent transactions committed at the source are immediately moved and delivered to the target.

Deferred Apply

For maximum flexibility, the Delivery module can apply data immediately or at a deferred time chosen by the user, without losing transactional integrity. This allows an additional layer of data protection when needed and keeps the secondary system at a consistent state behind the primary system. In this configuration, Oracle GoldenGate routes the changed data to the Trail File on the target server but does not deliver it to the target database until all captured changes have been delivered to the target Trail File. All changed data in the target Trail File can then be immediately applied to the target database, to bring it to a consistent state relative to the source, whenever the user chooses.

Oracle GoldenGate Manager

To give users control over Oracle GoldenGate processes, the Manager module 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 the Capture and Delivery modules
- Critical, informational event, and threshold reporting
- Resource management
- Trail File management

The Manager module executes requests on demand as well as unattended. For example, it can be used to restart Oracle GoldenGate components as well as monitor latency. The command-line interface lets users perform a variety of administrative, housekeeping, and reporting tasks. The module 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. Oracle GoldenGate offers transaction tracing flexibility to easily identify bottlenecks and tune the Oracle GoldenGate implementation for optimum performance.

For enhanced management of Oracle GoldenGate processes and solutions, customers should consider adding the Oracle Management Pack for Oracle GoldenGate.

Associated Products

There are two primary products that augment Oracle GoldenGate to enhance your real-time information platform:

- [Management Pack for Oracle GoldenGate](#). Management Pack for Oracle GoldenGate provides components that enable designing, configuring, managing, and monitoring Oracle GoldenGate core replication processes implemented across the enterprise.

- [Oracle GoldenGate Veridata](#). A data comparison utility that quickly compares data between two online databases and reports any discrepancies (can run as a standalone product). It supports heterogeneous databases including Oracle Database, HP NonStop, and Teradata.

Oracle Management Pack for Oracle GoldenGate

Oracle Management Pack for Oracle GoldenGate is a centralized, server-based graphical enterprise application that offers an intuitive way to define, configure, manage, and monitor Oracle GoldenGate processes. It leverages the management services of the core Oracle GoldenGate platform to help users reduce the deployment time for their continuous availability and real-time data integration configurations. This pack includes two robust and complementary products: Oracle GoldenGate Monitor, for monitoring Oracle GoldenGate implementations, and Oracle GoldenGate Director, for managing Oracle GoldenGate components. In addition, beginning with Oracle GoldenGate 11g Release 2, this pack includes a Plug-in for integrating Oracle GoldenGate Monitor with Oracle Enterprise Manager.

Oracle GoldenGate Monitor provides an end-to-end view of all of a user's Oracle GoldenGate implementations. GoldenGate Monitor can integrate with the existing alerting infrastructure with its sophisticated alerts based on SNMP, email and command line integration (CLI) for third-party call outs. GoldenGate Monitor is designed to take advantage of the low-impact agent that is included with Oracle GoldenGate, and supports Oracle GoldenGate Release 11.1.1.1.1 and higher. One agent per GoldenGate instance collects monitoring points—such as status, lag, number of inserts, updates, and deletes—and sends them to the GoldenGate Monitor Server. Using this data, Oracle GoldenGate Monitor graphically displays a customer's end-to-end replication solutions and provides sophisticated server side alerting logic.

Oracle GoldenGate Director is a graphical tool that expedites the Oracle GoldenGate process design and simplifies the management of GoldenGate replication components. GUI-based configuration reduces human errors and enables rapid production of multi-server solutions from a central location.

Oracle GoldenGate Veridata

Oracle GoldenGate Veridata is a high-speed, low-impact data comparison solution that identifies and reports data discrepancies between two databases, without interrupting those systems or the business processes they support. A standalone product, Oracle GoldenGate Veridata does not depend on the presence of Oracle GoldenGate's core components.

Using this application, companies can audit and verify large volumes of data shared across different business applications with absolute certainty. Oracle GoldenGate Veridata reduces the amount of time and the number of resources required to compare data, minimizes the impact of human errors, and accelerates the discovery of potential problems for rapid resolution.

One Platform, Many Solutions

Oracle GoldenGate easily supports a wide variety of topologies. These include one-to-one, one-to-many, many-to-one, and many-to-many—for both unidirectional and bidirectional configurations.

For unlimited scalability, cascading topologies can be created to eliminate any potential bottlenecks. By staging specific sets of database changes on the source or target system, different requirements can be met through a single pass on the data source. Each set of staged data can contain unique or overlapping sets of data.




Figure 2. Oracle GoldenGate supports numerous data propagation solutions to support real-time visibility across the enterprise.

Oracle GoldenGate can also integrate with JMS-based messaging systems for increased flexibility in distributing real-time data. This capability, along with the flat file delivery feature, allows Oracle GoldenGate to provide different architectures to augment existing investments.

This single real-time data movement platform enables companies to easily and successfully implement a variety of solutions for improving the availability, accessibility, performance, and integration of critical data across the enterprise. The most commonly implemented use cases for the Oracle Database include:

- Zero Downtime Migrations and Upgrades
- Query Offloading
- Disaster Recovery and Data Protection
- Active-Active Database Replication for Continuous Availability
- Operational Reporting and Real-time Data Warehousing

- Data Distribution and Synchronization for OLTP Systems

Zero Downtime Migrations and Upgrades

With its real-time heterogeneous replication capabilities, Oracle GoldenGate enables zero-downtime database, hardware, and OS migration. Through real-time, bidirectional data movement and synchronization between old and new systems, Oracle GoldenGate enables switchover from the existing system to the new system, when the new system is ready—without ever denying access to the application. The solution can support heterogeneous environments for cross-platform upgrades and database migrations from non-Oracle databases to Oracle databases including Oracle Exadata. In addition, the solution provides robust failback contingencies by keeping the old and new environments in sync in real time. By simultaneously using Oracle GoldenGate Veridata, users can identify and report data discrepancies across systems before switchover, without impacting the production environment.

In addition, Oracle GoldenGate is certified to offer zero downtime application upgrades for Oracle’s Siebel CRM and JD Edwards applications. Oracle GoldenGate also provides zero-downtime rolling upgrades for Oracle Database 10g or 11g to Oracle Database 12c.




Figure 3. Oracle GoldenGate enables organizations to upgrade and migrate mission-critical applications with no downtime.

Query Offloading

Oracle GoldenGate’s real-time data integration capabilities enable users to offload queries and read-only activity from production databases to a dedicated reporting server. With its support for heterogeneous systems GoldenGate can enable customers to offload OLTP systems running on mainframe systems, such as IBM DB2 and HP NonStop, as well as open systems Microsoft SQL Server and Sybase ASE, to Oracle Database or MySQL. This solution maximizes performance of the production system, reduces cost of ownership, and enables users to optimize the secondary reporting database for read-only activities.

Disaster Recovery and Data Protection

Oracle GoldenGate complements Oracle Active Data Guard's robust disaster recovery and high availability offering for the Oracle Database by providing the ability to use different Oracle Database versions, and different database or operating systems for the standby instances. In addition, Oracle GoldenGate supports Oracle Database Standard Edition and offers bidirectional data replication to run multi-master synchronization between Oracle Database or non-Oracle databases for continuous availability.

When configured for disaster recovery and data protection, Oracle GoldenGate provides a continuous availability solution that significantly improves recovery time for mission-critical systems. Oracle GoldenGate delivers up-to-the-second data to the backup system and enables immediate switchover to the new system if an outage occurs. It also immediately initiates real-time data capture from the standby database to update the primary system, once it is online, with any new data processed by the standby system.

Active-Active Database Replication

Oracle GoldenGate enables the highest availability and performance on critical systems by keeping two or more active databases in-sync as they support ongoing transactions. The product provides loop detection and built-in conflict detection/resolution capabilities to provide support for real-time data movement between active databases. This solution provides continuous availability for supported databases, better use of existing server assets, and improved system performance and scalability.

Operational Reporting and Real-Time Data Warehousing

Many operational reporting activities ideally seek to use the latest data available, and that often means running reports against production databases, which degrades performance. With Oracle GoldenGate, a secondary, cost-effective system can be deployed to serve the purposes of real-time, operational reporting, freeing up the resources of the critical source systems. Oracle GoldenGate 12c is certified to support operational reporting solutions for major Oracle applications including Oracle E-Business Suite, JD Edwards, PeopleSoft, Siebel CRM, Oracle Business Intelligence, Oracle ATG Web Commerce, and Oracle Coherence.




Figure 4. By continuously moving source data to a reporting database, users can access real-time information with no production impact.

Oracle GoldenGate can also be used for feeding data warehousing solutions with timely data. GoldenGate complements Oracle Data Integrator (ODI) by providing log-based data capture capabilities. The combined solution can be used to capture real-time data from source systems without performance impact and loading data warehouse tables with timely information. Oracle Data Integrator enables Extract, Load, Transform (E-LT) architecture that further improves performance and reduces cost of ownership by eliminating the middle-tier transformation server. Oracle GoldenGate 12c is even more tightly integrated with Oracle Data Integrator 12c. For example, Oracle GoldenGate Gate source and target systems are now configured as data servers in Oracle Data Integrator's Topology. Oracle GoldenGate parameter files can be automatically deployed to source and target Oracle GoldenGate instances through the JAgent technology, and Oracle GoldenGate instances can be started or stopped from Oracle Data Integrator. Both Oracle GoldenGate and Oracle Data Integrator can be used with Oracle Exadata as well as Oracle Database Appliance.

Data Distribution and Synchronization for OLTP Systems

Due to its multi-directional, real-time data movement capabilities, Oracle GoldenGate enables data distribution and synchronization for distributed systems. Its multi-directional replication configuration keeps data centers in synch for access to timely, accurate data across enterprise systems. Since Oracle GoldenGate's replication does not have distance limitations, typically customers implement Oracle GoldenGate for data distribution between data centers located in different regions and continents. It also enables continuous availability if one data center is not available. The product's heterogeneity enables these distributed systems to be different databases and OS platforms.

Oracle GoldenGate for Oracle Database

For the Oracle Database, Oracle GoldenGate can perform near real-time capture, transformation, routing, and apply of committed database DML transactions being written to redo and/or archive logs. Advanced options of the product provide for the capture, routing, and apply of Oracle-to-Oracle DDL operations, and the capture of DML transactions from Archive logs only.

Oracle GoldenGate can efficiently operate on any platform running the database; including Oracle Real Application Clusters (RAC), Oracle Exadata, and databases managed by Automatic Storage Management (ASM). The following sections describe GoldenGate functionality, while highlighting capabilities that are specific to Oracle GoldenGate for the Oracle Database.

Capture (Extract)

While the Oracle GoldenGate Capture process is running, it continually scans the redo logs looking for qualified transactions. GoldenGate runs as a database user and must have the appropriate database privileges to access the redo and archive logs. The Oracle Recovery Manager (RMAN) automatically retains the archive logs that are needed by Capture. In addition to supporting replication of DML, in an active-active Oracle environment, GoldenGate can also capture and apply DDL changes, including CREATE, ALTER and DROP operations on nearly all database objects, PL/SQL, and user privileges.

When a qualified transaction starts, the Capture process begins queuing the transaction details to memory. If the transaction details exceed the allocated process memory, the transaction details are written to temporary disk storage. This Capture mechanism enables GoldenGate to track and capture multiple transactions simultaneously. Additionally, the Capture module can track and capture transactions that span multiple redo or archive logs. When the transaction is committed to the database, the Capture process flushes all of the transaction details to a trail file for routing to the target database server(s). The data in this trail file is in commit order sequence.

Integrated Capture

Beginning with Oracle GoldenGate 11g Release 2, customers running Capture on an Oracle Database can choose to take advantage of the Integrated Capture capabilities. This multi-threaded Capture mechanism enables a variety of enhancements for customers extracting data from an Oracle Database. When operating in Integrated Capture mode, the Capture process integrates with an Oracle log mining server to receive change data from the server in the form of Logical Change Records (LCRs), which are a logical representation of the changed data. The Integrated Capture mechanism can seamlessly traverse between reading the in memory log buffer, the redo logs, and the archive logs, as necessary.

When using Integrated Capture, customers can choose to locate the Capture mechanism at either the source database or an alternate location that is of the same platform type as the source.

Performing the capture downstream is useful when it is necessary to minimize overhead at the source database. For example, when replicating changes from a mission-critical OLTP system to a reporting or test environment, it may be preferable to offload the capture process to the target location. Downstream capture can also be used to enable mining of earlier versions of the database in either real-time or archive log mode. The database where the downstream capture is being performed must be Oracle Database 11g Release 11.2.0.3 or higher; however, the logs may be from an Oracle Database 10g Release 2 or higher database.

Because it is fully integrated with the database, no additional setup is necessary to support Oracle Real Application Clusters (RAC), Automated Storage Management (ASM) or Transparent Data Encryption (TDE). Integrated capture handles point-in-time recovery and RAC integration more efficiently, automatically handling node outages.

Classic Capture

Oracle GoldenGate continues to support the existing Capture module, now referred to as Classic Capture, which directly accesses the database redo logs looking for DML changes to capture for distribution. It is possible to capture from redo logs stored inside of ASM. Adjusting the read size can improve Capture performance. In this mode, Capture can be integrated with Oracle RMAN to manage log retention.

To minimize overhead on the source database, customers can choose to use Oracle GoldenGate Archive Log Only (ALO) capture, which supports parsing the archive logs on a remote system. This system does not need to be identical to the target, for example, it is possible to parse Linux x86 redo logs on a Windows x86 server, or HP-UX PA-RISC logs on a Solaris SPARC system. When using ALO in connected mode, there are no data type or feature limitations. Using disconnected mode eliminates overhead on the source system.

Routing (Data Pump)

A separate Capture process continually scans the staging trail file, awaiting new data. When new data is detected in the staging trail file, it is packaged for routing via TCP/IP to specific target locations. The target location can be a single server disk location, multiple disk locations, or multiple servers and disk locations.

Oracle GoldenGate uses TCP/IP, including IPV6, for sending data, so no geographical distance constraints are imposed between the source and target systems. Advanced options provide for encryption (using FIPS or Blowfish) and compression of the data within the TCP/IP packet. At the target locations, a communications process receives the incoming transmission from TCP/IP, decrypts and decompresses the data packet, and writes the transaction information to a local trail file. An acknowledgement is returned to the Data Pump that the data was received and written to the trail file successfully.

“Oracle GoldenGate 12c’s new Integrated Delivery mode for Oracle Database auto-tunes the parallelism and computes dependencies between transactions automatically, which allows end users to easily configure the product with a single Delivery process. This advanced replication, tested with an LHC application workload in the context of the CERN openlab collaboration with Oracle, performs five times faster compared to previous GoldenGate version and simplifies the management of a multi-tier replication environment.”

--Eva Dafonte Perez, IT, CERN“

Delivery (Replicat)

At the target database locations, Oracle GoldenGate also runs as a database user. Database privileges provide access that allows the Delivery process to execute DML and DDL on designated schemas and tables. The Delivery process continually scans the trail file for incoming data. When new data is detected, the source database transaction DML or DDL statement is recreated and executed.

In an active-active configuration, conflicting updates can be detected using Oracle GoldenGate’s built-in conflict detection mechanism. Once detected, these conflicts can be automatically resolved using one or more of the supplied conflict resolution methods or by a custom method, as needed to meet business requirements.

Integrated Delivery

Customers delivering data to an Oracle Database 11g Release 11.2.0.4 or Oracle Database 12c Release 12.1.0.1 and higher database, can improve performance and provide better scalability and load balancing by using Integrated Delivery. The new process simplifies configuration and increases performance more than 100 percent. Minimal changes are required to implement this change, which leverages the database parallel apply servers for automatic dependency aware parallel apply. With Integrated Delivery, there is no need for users to manually split the delivery process into multiple threads and manage multiple parameter files.

The Oracle GoldenGate Delivery process continues to read the Trail File, just as it did before, however, with Integrated Delivery, it constructs Logical Change Records (LCRs) that are used to describe the changed data. The Delivery process then transmits these LCRs to the Inbound Server of the Oracle Database using a lightweight streaming API.

The Inbound Server consists of the Receiver, which reads the LCRs, along with the Preparer, Coordinator, and one or more Appliers. The Preparer computes dependencies between the transactions and then groups and sorts them in dependency order. Next, the Coordinator maintains the order between the transaction groups as they are handed off to multiple Appliers. The Database Apply processes perform all necessary changes for the assigned transactions, including conflict detection and error handling.




Figure 5. Oracle GoldenGate 12c Integrated Delivery for the Oracle Database

Supplemental logging of scheduling columns (Primary Key, Unique Index and Foreign Key) are needed to support the dependency aware apply processing at the target. Integrated Delivery can process non-Oracle database trail files as long as all of the scheduling columns are logged by the non-Oracle database and are available in the trail file. The Integrated Delivery feature also provides support for the Oracle Multitenant Database feature at the Pluggable Database level. Just as shown in the prior section, multiple Integrated Delivery processes can read the same trail file and apply changes to multiple PDBs at the target.

Users running on an Oracle Database version prior to 11.2.0.4 can still continue to use the Classic Delivery process in their Oracle GoldenGate deployments. However, Oracle GoldenGate recommends that users upgrade to Oracle Database 11.2.0.4 or higher to take advantage of the new Integrated Delivery process.

Coordinated Delivery

Customers delivering data to non-Oracle data stores (or Oracle Database versions before 11.2.0.4) who find it necessary to split their delivery process into multiple threads can use the Coordinated Delivery feature available with Oracle GoldenGate 12c to eliminate the need to manage multiple parameter files. In addition to requiring a single parameter file for multiple Delivery processes, Coordinated Delivery also automatically provides coordination across selected events that require ordering, including DDL, Primary Key updates, Event Marker Interface (EMI) and SQLEXEC. Coordinated Delivery can be used with both Oracle and non-Oracle data stores.




Figure 6. Oracle GoldenGate 12c Coordinated Delivery

Transformation

Data transformation can occur within any step of the capture, route, and delivery process; however, it is performed most commonly during the data delivery phase. Oracle GoldenGate transformation supports capture and delivery between dissimilar databases, between schemas and tables with different names, between source and target columns with different names or data types, as well as row-level transformations, where data can be modified based upon user requirements.

Monitoring and Managing

Oracle provides a variety of tools for monitoring and managing Oracle GoldenGate. Beginning with Oracle GoldenGate 11g Release 2, Oracle Management Pack for Oracle GoldenGate includes a plug-in for integrating Oracle GoldenGate Monitor with Oracle Enterprise Manager. The plug-in for Oracle Enterprise Manager Cloud Control lets users gather configuration data and track configuration changes for Oracle GoldenGate instances. It also provides users with the capability to raise alerts and violations based on thresholds set on monitored targets and configuration data.

Oracle Database 12c Automatic Workload Repository (AWR) now includes replication statistics for Oracle GoldenGate. The Automatic Workload Repository is a repository of historical performance data, such as system resource usage, capture rate and apply coordinator rate, that can be helpful when you are making performance tuning decisions for your replication environment.

Additionally, a new package, `dbms_goldengate_auth.grant_admin_privilege`, greatly simplifies privilege management by allowing all of the privileges needed to administer a replication environment to be granted with a single stored procedure call.

Finally, there are new views available to further simplify management of a replicated database.

Optimized for Oracle Database

Oracle GoldenGate 12c has been optimized to support a variety of new functionality available with Oracle Database 12c.

Integration with Oracle Data Guard

Oracle GoldenGate with Integrated Capture is designed to interoperate with Oracle Data Guard, to provide the highest levels of database scalability and availability. Should the primary site become unavailable, it is possible for Oracle GoldenGate to switch to using the physical standby or Active Data Guard standby database as the primary site. Users can restart the Oracle GoldenGate environment at the new primary site.

Integration with Oracle Data Guard FSFO provides automated and transparent failover of Oracle GoldenGate components with the failover/switchover of the primary database so replication can continue without any manual intervention. In the instance of a failover/switchover, Oracle GoldenGate processes are started on the standby system during the role transition and insure that no data is lost during the failover/switchover.

GoldenGate Downstream Capture is able to utilize the Data Guard redo log transport mechanism to process transactions off-source. There is no chance of data loss when using the Data Guard Redo log transport mechanism to ship redo logs to a downstream / remote machine where GoldenGate Capture processes are installed.

Support for Oracle Multitenant Container Database Feature

Oracle GoldenGate 12c also supports the new Oracle Multitenant Container Database (CDB), which allows a common pool of resources to handle multiple databases (known as containers). Containers with user data are called Pluggable Databases (PDBs). These PDBs can be unplugged from one CDB and plugged into another CDB.

Oracle GoldenGate customers can mine multiple PDBs with a single Integrated Capture process and can write changes from multiple PDBs into a single trail file. Capturing changed data at the container level rather than at the individual databases reduces overhead. Memory and processes are required at the CDB level only. In Oracle Database 12c, all of the PDBs in a CDB must have the same attributes, such as character set, locale, and case-sensitivity. The Integrated Capture process must be configured to run at the root container as a common user in order to interact with the log mining server.

The Delivery process can connect and apply changes to a single pluggable database. To configure replication from multiple source pluggable databases to multiple target pluggable databases, users can configure parallel Integrated Capture and Delivery streams, each handling data for one pluggable database. Alternatively, users can configure a single Integrated Capture that captures from multiple source pluggable databases and writes to one trail that is read by multiple Delivery processes, each applying to a different target pluggable database. Yet another alternative is to use one Integrated Capture writing to multiple trails, each trail read by a Delivery process assigned to a specific target pluggable database as shown below.




Figure 7. Oracle GoldenGate 12c Integrated Capture and Delivery for Oracle Multitenant Container Database

Support for Cloud Environments

As organizations and individual departments adopt cloud applications, in many cases they create point-to-point connections that bypass well-established integration principles. Lack of centralized monitoring and management causes extra work for system administrators, especially as cloud data and functions are shared with on-premise information systems. Data integration solutions play a vital role in migrating data simply, efficiently, and reliably to the cloud. They are also essential to platform-as-a-service (PaaS) implementations because they support cloud deployments with data-layer application integration between on-premise and cloud environments.

Consolidating into Private Cloud

Oracle Data Integrator and Oracle GoldenGate can connect on-premise enterprise systems to a private cloud by moving data in bulk or as real-time transactions across geographies. GoldenGate's real-time and bidirectional data replication capabilities between heterogeneous systems enable companies to consolidate data without interrupting business operations. The new Oracle Database 12c release introduced new multitenant architecture, which simplifies the process of consolidating databases onto the cloud, enabling customers to manage many databases as one - without changing their applications. Oracle GoldenGate 12c is optimized for Oracle Database 12c and supports its new multitenant architecture, making it a perfect solution for consolidation without impacting business operations. A failback option allows the IT team to test the target environment as long as necessary, minimizing risk.

Integrating Public Cloud Environments with On Premises Systems

The trend for enterprises to utilize both cloud-based and on-premises systems to host their database instances is increasing. This trend brings the importance of implementing reliable real-time replication across the cloud and on-premises environments to make sure data remains fresh

and relevant for business operations. Some of the key requirements in keeping the data fresh in hybrid environments include – minimal latency, reliability and security. Data must remain fresh. As data ages it becomes less relevant and less valuable—day-old data is often insufficient in today’s competitive landscape. Reliability must be guaranteed despite system or connectivity issues that can occur between the cloud and on-premises instances.

Oracle GoldenGate offers real-time data replication between on premises databases and databases hosted on public cloud environments to enable reliable, secure and timely data integration. For more information please review the knowledge article (ID - 1588484.1) titled ['Replicating between Cloud and On-Premises using Oracle GoldenGate'](#).

Support for Additional Oracle Database Functionality

In addition to CDB support, with Integrated Capture, Oracle GoldenGate provides support for compression, including Basic Compression, OLTP-compression, and Exadata Hybrid Columnar Compression (EHCC) as well as distributed transaction support for XA and PDML transactions on Oracle RAC. XML Object Relational (OR) and XML Binary data types support has also been added along with LOB full and partial read (selective update) from the redo log. Integrated Capture also extends the supported VARCHAR2, RAW and NVARCHAR2 limit from 4000 to 32764.

Customer Examples

As described earlier, there are a variety of different ways that customers can choose to configure a replication solution using Oracle GoldenGate. Here are just a few of the many real-life examples of how customers have chosen to improve performance and enhance availability by implementing Oracle GoldenGate in an Oracle Database environment.

1) Credishop Shortens Credit Card Transaction Time by 80% through Database Replication

Credishop—administrator for its own brand of credit cards and for those of other companies and partners—decided to improve its IT infrastructure to sustain its rapid growth by implementing Oracle GoldenGate. The company used the solution to replicate its databases, creating identical databases so that each one can attend to distinct business needs. One database is dedicated solely to processing credit card transactions (which is the company’s most critical operation), and the other handles routine operations, such as call center service and billing.

Credishop ensured high availability and high processing speed for both environments, reducing the time needed to process credit card transactions by 80%, which translates to faster customer service and improved productivity. Credishop’s Oracle Database environment supports more than 900 simultaneous users and handles approximately 800,000 transactions per month without interruptions or delays.

“Oracle GoldenGate enabled us to replicate all of the most important and critical data in our Oracle Database, Enterprise Edition, which handles 95% of all data processed by Credishop, including credit card transactions. With this combination, we realized exceptional performance gains and can now authorize a sale in just three seconds said Edgar Galan Rodrigues, Database Administrator, Credishop.

“We began using Oracle when we were founded in 1999, and we trust Oracle’s solutions 100%. Our adoption of Oracle GoldenGate 11g brought a series of gains to our business. With greater speed and agility, we have greatly improved customer service.”

— Edgar Galan Rodrigues , Credishop

2) Avea İletişim Hizmetleri A.Ş. Optimizes Database Compression, Reporting, and Campaign Analyses

Avea İletişim Hizmetleri A.Ş., the sole GSM 1800 mobile operator of Turkey, provides services to 97% of Turkey's population through its next generation network. With roaming agreements covering 199 countries, Avea is growing rapidly and currently serves 12.8 million subscribers.

One of their top challenges was to create an open telecommunications platform that can integrate new services such as mobile number portability and 3G network services. The company wanted to enable fast reporting and immediate access to business data across the company by establishing an efficient and flexible data warehouse. The ultimate goal was improving marketing campaigns by creating more specific target audiences and running more detailed post-campaign analyses.

Avea migrated data from disparate systems with Oracle GoldenGate and created a central repository on Oracle Database, raising data access performance by a factor of five. The company used the advanced data compression capabilities of Oracle Exadata Database Machine and Oracle Database to decrease 40 terabytes of data to a mere 10 terabytes.

Avea leveraged Oracle Data Integrator to optimize its Extract-Transform-Load (ETL) processes, decreasing the ETL batch window by more than 30%. Oracle’s solution accelerated the company’s reporting performance by three to ten times by integrating platforms, source systems, and operational data stores with Oracle GoldenGate.

Avea also used the advanced visualization and analysis features of Oracle Business Intelligence Enterprise Edition, enabling more efficient and in-depth analysis of data from across the company’s telecommunications operations. Using Oracle’s solution Avea automated most data warehouse maintenance processes and lowered the maintenance effort by nearly half.

3) Aria Systems Ensures Cloud Subscription and Billing System High Availability and Performance with Robust, Highly Secure Data Infrastructure

Aria Systems is a leading provider of cloud billing and subscription management solutions for Global 2000 companies. The Aria Subscription Billing Platform manages the entire subscriber lifecycle for all recurring revenue models as a company grows and expands its business. Pitney Bowes, Ingersoll Rand, DreamWorks, EMC, Internap, VMware, and Hootsuite all rely on Aria Systems for fast time-to-market, low operational costs, and operational flexibility.

Aria Systems looked to deliver an end-to-end cloud billing and subscription management solution that was affordable and enabled rapid onboarding for customers. Furthermore, it needed to ensure 100% availability of its subscription billing platform, on which customers rely to capture revenue streams from the online services they offer. The company selected Oracle Database, Enterprise Edition; Oracle GoldenGate; and Oracle Real Application Clusters as the foundation for its solution. Oracle has enabled Aria Systems to deliver the high performance and availability that its customers expect, as well as the scalability to support rapid spikes in volume, such as accommodating a ten-fold transaction increase over the holidays without performance degradation. Aria Systems also ensured zero downtime for scheduled maintenance by using the live-release capabilities in Oracle Database

“Oracle Database, Enterprise Edition 11g; Oracle GoldenGate; and Oracle Real Application Clusters provided us with a very secure and stable environment to support our end-to-end cloud billing and subscription management solution. Oracle’s solutions are among the most critical components of our infrastructure—ensuring 24/7 reliability and availability with superior fault tolerance.”

— Oleg Ganopolskiy, Aria Systems

Key Benefits of Using GoldenGate for Oracle Database

As these customer examples attest, Oracle GoldenGate brings various IT and business benefits to Oracle Database users. Here are the main areas.

- **Cut costs and improve efficiencies:** Oracle GoldenGate’s high-performance, bidirectional data replication capabilities enable companies to eliminate downtime for unplanned and planned outages, which can be cost prohibitive. The low-overhead nature of the solution enables customers to better use their existing hardware and network investments. In addition, it allows companies to offload resource intensive

operations to a lower-cost replica of the source system to lengthen the lifetime of existing systems.

- **Reduce risk, ensure business continuity:** Oracle GoldenGate minimizes data loss and downtime risks during migration, upgrade, and maintenance activities by offering a fallback option to the old environment. In addition, it moves data with extreme reliability and transactional integrity. By eliminating downtime during both planned and unplanned outages GoldenGate helps companies achieve business continuity, especially in high-risk operations such as in healthcare and financial services.
- **Improve business operations.** Oracle GoldenGate's zero downtime system migration solution allows companies to move to a better infrastructure sooner rather than later, since GoldenGate enables uninterrupted operations during such projects. Businesses can enhance service levels and decrease costs, and achieve other productivity benefits earlier. Maintaining SLAs for availability in general enables customer satisfaction and competitive advantage in a world where end users demand 24/7 availability. In addition, GoldenGate delivers timely information for BI and reporting activities. By leveraging up-to-date data companies can have better business insight to provide an enhanced customer experience and achieve greater operational efficiencies.

Oracle GoldenGate delivers these business and IT benefits with a very easy-to-use platform that reduces total cost of ownership and delivers high performance, flexibility, and reliable data integration and replication across heterogeneous systems.

Oracle GoldenGate is tightly integrated with Oracle's technology stack especially Oracle Database. Compared to 3rd party vendors, GoldenGate offers the most optimized and flexible solutions for maximum availability and real-time data integration for Oracle Database users.

Summary

With Oracle GoldenGate's real-time, heterogeneous database synchronization capabilities, Oracle Database users can replicate both DML and DDL operations. The product allows users to move, filter, and transform data between similar or dissimilar supported Oracle Database versions, or move, filter, and transform data between an Oracle Database and a database from another vendor.

Oracle GoldenGate enables the world's largest enterprises to improve the availability, performance, and accessibility of the transactional data that drives mission-critical business processes. Oracle GoldenGate's wide variety of use cases includes real-time business intelligence, query offloading, zero-downtime upgrades and migrations, and active-active database replication for data distribution and, continuous availability.



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

ORACLE®

Using Oracle GoldenGate for Oracle Database
September 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

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.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. 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. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0410