

ORACLE

Oracle GoldenGate Innovations

Another 30 Years of Trusted Real-time Data



Executive Summary

Enterprise AI has moved past the question of model access and quality. Most organizations can easily procure advanced frontier models, copilots, and AI development tools. The question that now separates leaders from laggards is whether those systems can operate on current, governed, trustworthy business data across the operational reality of the enterprise. This is precisely where most AI initiatives stall. The bottleneck is no longer intelligence. It is data: fresh, ordered, trusted, and available everywhere it is needed, in real-time.

This is the problem Oracle GoldenGate was built to solve, long before generative AI made it urgent. For more than thirty years, GoldenGate has been the enterprise standard for real-time, transactionally consistent data movement, used in production by approximately 80% of the Fortune 500 and nearly all major banking and financial institutions [5]. It operates with extreme reliability and trustworthiness, with exactly-once guarantees and sub-second latency.

This paper discusses two main themes for different audiences. The first is a Business Case for Enterprise Leadership. This is meant to cover discussion topics that a technical leader can carry to the CIO, CDO, and CFO: why real-time, governed data integration has become a board-level AI-readiness issue rather than a platform-hygiene project, and where the return on that investment shows up. The remainder of this paper is the technical substantiation: what makes GoldenGate structurally different, how it performs in the world's most demanding production environments, the architecture patterns it supports, and ten areas of innovation propelling the platform into its next thirty years.

Those ten innovations include:

- **Data Fabric & AI Data Platform** – unified, real-time data products for AI and analytics
- **Kernel Integration with Oracle AI Database** – performance no competitor can match
- **Multicloud Deployment Everywhere** – OCI, Azure, Google Cloud, AWS, and beyond
- **Autonomous & AI-Assisted Operations** – simplicity at mission-critical scale
- **Real-Time Data Fabric Architecture** – the strategic control plane for data movement
- **Containers & Kubernetes Everywhere** – developer-first, cloud-native deployment
- **Observability, Governance & Trust** – auditable, verifiable, exactly-once pipelines
- **GoldenGate for AI & Agentic Workloads** – embeddings, LLMs, and RAG pipelines
- **Fully Managed Multicloud Platform** – OCI GoldenGate with 99.95% SLA and auto-scale
- **Consumption-Based & Free Pricing** – GoldenGate Free and elastic subscriptions

What follows explores these innovations in detail and explains how they drive results for organizations embarking on AI transformation—where real-time, trusted data is the indispensable foundation.




Table of Contents

Why This Matters for Enterprise Leaders	5
Why Oracle GoldenGate	7
What Makes GoldenGate Different	8
Log-Based Capture vs. API-Based Polling.....	8
Transaction Semantics & Exactly-Once Guarantees	8
The Competitive Moat.....	9
GoldenGate in a Competitive Market.....	9
From CDC Tool to AI Data Platform	9
Distributed Mesh vs. Hub-and-Spoke	10
The Native Wire Protocol Advantage	11
Oracle Kernel Integration: A Structural Advantage	12
Fully Managed Multicloud: Not a Marketplace VM	13
30 Years of Innovation	14
Continuous Availability and Disaster Recovery	14
Trickle-Feed Analytic Data Warehouses	15
Streaming Pipelines and the AI Era	16
Everyday Moments by GoldenGate.....	16
Proof Points: The Numbers Behind the Story.....	17
Real-World Throughput.....	17
Customer Reference Points.....	18
When to Use GoldenGate.....	18
Common GoldenGate Architecture Patterns	19
Pattern 1: Operational High Availability.....	20
Pattern 2: CDC to any Kafka-compatible Topics	20
Pattern 3: Native, Pub-Sub Data Streams.....	21
Pattern 4: Real-Time Lakehouse Ingestion.....	21
Pattern 5: Real-Time AI Pipeline	22

Innovation for the Next 30 Years	22
1. Data Fabric & the Backbone for Enterprise AI Data	22
2. Innovation with Oracle AI Database.....	24
3. Multicloud Deployment Everywhere	25
4. Autonomous & AI-Assisted Operations	26
5. Data Mesh Service Topologies.....	27
6. Containers, Kubernetes & Developer-First Deployment	27
7. Security, Observability, Governance & Trust.....	28
8. GoldenGate for AI & Agentic Workloads	30
9. Simple, Fully Managed Multicloud Platform	31
10. Inexpensive, Consumption-Based & Free Pricing	31
Reference Architecture: A Real-World Pipeline	32
The Scenario: Real-Time Financial Intelligence	32
Latency Expectations	34
Failure Handling and Recovery.....	34
New Frontiers	35
Decoupling your Data from the Lakehouse.....	35
An Open and Independent ZeroETL Layer.....	36
GoldenGate goes Interplanetary	37
The Next 30 Years: GoldenGate stays Gold	38
Limitations & Tradeoffs.....	38
Conclusion	39
Appendices	41
Glossary of Terms	41
FAQ: Oracle GoldenGate	44
Citations and References.....	46

Why This Matters for Enterprise Leaders

Access to AI models is no longer the constraint. Any enterprise can buy advanced LLMs, copilots, and development tools. The question that decides whether AI pays off is whether those systems can operate inside the real enterprise—across legacy databases, ERP systems, SaaS applications, lakehouses, APIs, and regulated workflows. That is where most initiatives stall. MIT reports that 95% of organizations see no measurable P&L impact from generative AI despite \$30–\$40 billion in spending [1]; RAND finds that more than 80% of enterprise AI projects fail, with roughly a third abandoned before they reach production [2].




The pattern is consistent: AI projects do not fail because the model is bad or AI features are missing. AI projects most often fail because the enterprise cannot reliably connect AI to current, governed, trustworthy business data. 82% of IT leaders cite data integration as a top challenge, with teams spending more than a third of their time building custom integrations [3]. Gartner predicts that through 2026, organizations will abandon 60% of AI projects that are not supported by AI-ready data [4]. For the leaders accountable for that data—CIOs, CDOs, CDAOs, CTOs, and infrastructure executives—this reframes data integration as a board-level AI-readiness issue, not a platform-hygiene project.

For organizations built on enterprise applications and operational databases, their AI strategy now depends on that very same data, sourced from all those existing systems. Oracle GoldenGate matters in the AI era because it delivers the one thing AI and Agentic systems cannot function without: *reliable, trusted data truth from an entire ecosystem of operational systems.*

AI systems, and agents above all, are only as useful as the operational reality they can see. Analytical AI can sometimes tolerate yesterday's batch load. Operational AI cannot. An agent acting on stale inventory, an outdated account status, an expired permission, or an incomplete customer history does not become cautious, *it becomes confidently wrong.* GoldenGate streams committed business events as they happen, preserving transaction order, integrity, and consistency. That gives AI models, RAG systems, vector databases, lakehouses, and event-driven applications access to current data reality rather than a delayed snapshot of it.

In a simple AI pilot project, an AI assistant answers questions with limited consequences. In a production system, an AI agent triggers workflows, updates records, generates customer responses, routes exceptions, and initiates downstream business actions with real fiscal consequences.

Bottom-line insight: AI ROI is constrained by whether enterprises can connect, govern, observe, and operationalize trusted data across fragmented systems.



That shift raises the stakes on four capabilities at once: the data must be current, access to it must be governed, every action must be auditable, and the system as a whole must be observable enough to control. Leadership must be able to answer what data an agent saw, whether that data was current, what action the agent took, what system it touched, and what the action cost.


For decision-makers, the return on that investment is concrete and falls into four areas:

- Trusted, real-time data. CDC, streaming, and governed integration let AI act on current operational reality, reducing lost-trust incidents and improving production reliability—the difference between an agent that is right and one that is confidently wrong.
- Engineering productivity. Reusable connectors and data pipelines replace custom integration code, accelerating delivery and redirecting scarce engineering capacity from integration toil to value creation.
- Governed access and risk reduction. Policy, identity, and runtime governance ensure agents reach only the data and systems they are permitted to use—the technical enforcement layer beneath every AI governance and compliance commitment.
- Observability and cost control. Per-agent and per-pipeline observability shows leaders what each initiative costs, which ones are working, and where funding should be redirected.

In the AI era, GoldenGate is therefore not simply a replication tool, **it is a strategic control point for trusted, real-time enterprise data**. GoldenGate is the layer that makes the existing IT investments (e.g. DBs, cloud platforms, analytics, and AI tools) work together rather than in isolation. For any enterprise pursuing more cost-effective, efficient and autonomous operations, GoldenGate provides the foundation to move from experimentation to production value from AI/Agentic initiatives.

Why Oracle GoldenGate

AI initiatives don't stall from lack of intelligence; *they stall from data trust issues*. Answering that challenge requires a platform engineered for exactly the properties AI needs: currency, order, trust, and reach across every system the enterprise runs. That platform is Oracle GoldenGate.



The Oracle GoldenGate platform is designed to connect data consumers with data producers. It sets the enterprise standard for mission-critical, real-time data movement used in the applications that power everyday life—ATM cash machines, point-of-sale systems, eCommerce platforms, airline booking systems, social media, and mobile apps. With more than 8,000 customers globally across banking, finance, retail, telecommunications, manufacturing and government, GoldenGate is the most trusted and widely deployed data integration platform in the world [5].

Typical technical use cases include:

- Database High Availability and Zero Downtime Operations
- Real-time Data Ingest for Analytics and AI
- Streaming Pipelines for Lakehouses and Data Platforms
- AI-Ready Data Fabric and Embedding Generation
- Online Migrations and ZeroETL

For more than 30 years, GoldenGate has been the industry-leading data integration platform in the market. Originally a startup in San Francisco, GoldenGate Software specialized in disaster recovery and high availability for the Tandem NonStop database—the backbone for many banking platforms. Oracle acquired GoldenGate in 2009 and has continued to invest deeply in the platform across on-premises, cloud-native, and now AI-native capabilities. The result is a product that is more essential today than at any point in its history. Approximately 80% of the Fortune 500 run GoldenGate—and that number is growing as AI transformation makes real-time, trusted data the most strategic asset any enterprise can own.

What Makes GoldenGate Different

Every enterprise software market has its share of contenders. Data integration is no exception—there are open-source tools, cloud-native alternatives, and a steady stream of challengers claiming to match GoldenGate’s capabilities. After thirty years, GoldenGate is still at the top because of a set of fundamental technical advantages that are genuinely difficult to replicate. Understanding these differences is essential for any architect or decision-maker evaluating real-time data infrastructure.

Log-Based Capture vs. API-Based Polling

Most data integration tools move data in one of two ways: they either poll the source database with SQL queries at regular intervals, or they call application-layer APIs to extract changed records. Both approaches share a common flaw—they interact directly with the live system, creating CPU load, I/O pressure, and query contention on the very applications they are supposed to serve.

GoldenGate takes a fundamentally different approach. It reads directly from the database transaction log—the redo log in Oracle, the WAL in PostgreSQL, the binary log in MySQL—capturing every committed transaction at the moment it is written, before it ever touches application-visible tables. This is the meaning of “minimally invasive.” GoldenGate is an observer of what the database is already doing, not a participant in it. The source application sees no additional load from GoldenGate’s capture process, because GoldenGate is reading a stream that the database was already writing for its own recovery purposes.

This distinction matters enormously at scale. A global bank processing millions of payment transactions per hour cannot afford to have its replication tool competing with account holders for database resources. A manufacturing company running a just-in-time supply chain cannot tolerate the query overhead that API-based integrations impose on its transactional systems. Log-based capture is not just a technical nicety—it is the architectural choice that makes GoldenGate viable in the world’s most demanding environments.

Transaction Semantics & Exactly-Once Guarantees

Here is where GoldenGate separates itself from every alternative in a way that is not merely a performance advantage but a correctness guarantee. Most streaming and integration tools operate on an at-least-once delivery model: in the event of a failure, they replay events and accept that some events may be delivered more than once. For analytics workloads, a duplicate row is an inconvenience. For financial transaction processing, payment systems, or any mission-critical operational database, a duplicated transaction is a potential disaster.

GoldenGate provides exactly-once, transactionally consistent delivery. Every transaction arrives at the target exactly once, in commit order, with the same transactional boundaries as the source. This is made possible by GoldenGate’s checkpoint architecture—the platform maintains a precise position in



the source transaction log and in the target apply queue, allowing recovery from any failure point without re-delivering already-committed transactions.

What this means in practice is that you can trust GoldenGate's output as the authoritative record of what happened in the source database, in the order it happened, with nothing added and nothing missing. This is the property that makes GoldenGate the only replication platform certified for financial settlement, regulatory reporting, and Oracle Maximum Availability Architecture Platinum and Diamond tiers [6]. If your use case requires strict transaction ordering and zero duplicates across systems, GoldenGate is the correct choice—no alternative offers this guarantee at enterprise scale.

The Competitive Moat

Beyond log-based capture and exactly-once delivery, GoldenGate's competitive position rests on a set of advantages that have compounded over thirty years. The platform supports over 1,000 combinations of databases, operating systems, chipsets, and cloud environments—a breadth of coverage that no competitor comes close to matching. Its kernel-level integration with Oracle Database means that GoldenGate can access private APIs unavailable to any third party, achieving performance ceilings that are structurally unreachable for alternatives. Its Veridata module provides algorithmic proof that source and target data are identical—a capability with no equivalent in the market.

Perhaps most importantly, GoldenGate's production heritage in the world's most demanding environments—global payment networks, central bank clearing systems, tier-one telecommunications infrastructure—has produced a depth of edge-case handling and failure recovery logic that only comes from three decades of continuous operation under real-world conditions. Open-source tools and newer entrants are often excellent for typical workloads. They tend to encounter GoldenGate's domain when typical is not enough.

GoldenGate in a Competitive Market

The data integration market has never had more participants. Debezium, Qlik Replicate, AWS Database Migration Service, IBM IADR, Striim, HVR—each of these tools has earned a legitimate place in the ecosystem, and each has real strengths that merit respect. A fair competitive analysis acknowledges those strengths while being equally honest about where GoldenGate occupies a structurally different tier of capability. GoldenGate is designed for a different class of workload than most alternatives. Understanding why requires looking at the architecture beneath the feature lists.

From CDC Tool to AI Data Platform

The most important competitive distinction in 2026 is not about individual features—it is about platform trajectory. Tools like Debezium, Qlik Replicate, AWS DMS, and IBM IADR were all designed, at their core, to solve the CDC and replication problem: capture database changes and move them

somewhere else. They do this with varying degrees of sophistication and in various operational contexts, and for straightforward replication scenarios, several of them do it quite well.

GoldenGate started from a more demanding foundational problem thirty years ago, its focus was on high availability for financial systems. And since then the platform has not stood still. GoldenGate Stream Analytics 26ai adds stateful, real-time stream processing with LLM operators and AutoML directly within streaming pipelines. GoldenGate Data Streams introduces a native Pub/Sub API layer based on AsyncAPI, eliminating the need for a separate Kafka cluster in many event-driven architectures. Veridata provides algorithmic data verification and auditability for compliance-grade governance. OCI GoldenGate delivers a fully managed multicloud service operating under a 99.95% SLA. The GoldenGate AI Microservice integrates natively with OCI Generative AI, OpenAI, Azure AI Foundry, AWS Bedrock, Gemini, Claude and any LLM-compatible API.

	GoldenGate	Others...
Created for mission-critical HA scenarios	✓	✗
Built-in stream processing & pipelines	✓	✗
Pub-sub API for AI agents & app clients	✓	✗
Multicloud native, not just 'marketplace'	✓	✗
Multimodal Agentic, LLM & ML integration	✓	✗
Kernel integration to Oracle AI Database	✓	✗
100's of third-party integrations	✓	✗

None of GoldenGate's primary competitors offer anything approaching this breadth. Debezium, for all its community momentum and developer-friendliness, remains a CDC capture library that delegates distribution, transformation, analytics, governance, and AI integration entirely to other tools in the stack. Debezium is also very limited with regard to the source and target data stores it supports. Qlik Replicate and IBM IADR are established enterprise CDC platforms with strong database coverage and solid operational track records, but they are not as focused on mission-critical use cases and their product roadmaps have not pivoted toward AI-native data platforms at the same pace or depth. AWS DMS is a capable migration tool within the AWS ecosystem but is architecturally constrained for multicloud, mission-critical, or long-running continuous replication scenarios.

GoldenGate's competitors are still primarily answering the question of how to move data. GoldenGate is now answering the question of what data can become.


Distributed Mesh vs. Hub-and-Spoke

Architecture topology is one of the least-discussed but most consequential differences between GoldenGate and its competitors. Most CDC tools are built around a hub-and-spoke model: a central broker or Kafka cluster sits at the middle of the architecture, and all data flows in and out through that hub. This works acceptably when all of your data is in one place or when Kafka is already a funded,



managed part of your infrastructure. It becomes a significant liability when your data is distributed—across data centers, across cloud providers, across continents—because every spoke now has to route through a centralized choke point that adds latency, complexity, and cost.

GoldenGate's architecture is fundamentally different. It is a distributed mesh—a network of lightweight, independently deployable microservices that can be arranged in any topology the workload demands. A hub configuration is entirely possible when centralization makes sense. But GoldenGate can equally be deployed as a true mesh, with GoldenGate agents running at the edge of every source and target system, communicating peer-to-peer across the WAN, with no central bottleneck in the critical data path.



For the kinds of globally distributed organizations that depend on GoldenGate—multinational banks, global airlines, hyperscale technology companies—this topological flexibility is not a nice-to-have. It is the only architecture that can serve their requirements without compromise.

Debezium's model is particularly instructive here. It is an excellent capture library, but it has no native, fully-formed distribution layer at all—the vast majority of users produce to Kafka, and Kafka becomes your distribution infrastructure. This means that a Debezium-based architecture inherits Kafka's operational footprint, Kafka's throughput characteristics, and Kafka's latency floor across the WAN. For teams already invested in Confluent or a managed Kafka service, this may be an acceptable trade. For teams that need sub-second, exactly-once delivery across cloud boundaries without managing a Kafka cluster, it is a meaningful constraint.

The Native Wire Protocol Advantage

Here is a technical dimension of GoldenGate's design that rarely appears in analyst comparison matrices, but matters enormously in practice. Every major competitor—Debezium, Qlik, DMS, IIDR—relies on either JDBC/ODBC database connections or a third-party messaging layer such as Kafka or IBM MQ to move data between systems. This is a reasonable engineering choice, but it comes with real constraints: JDBC connections create load on the source database, external messaging layers add


latency and operational overhead, and neither approach was designed specifically for the demands of high-throughput, mission-critical replication across wide-area networks.

GoldenGate operates its own native binary wire protocol for transmitting data across the WAN. This is not a minor implementation detail—it is a purpose-built transport that was designed from scratch for exactly this workload. The protocol transmits compressed, binary-encoded transaction records over WebSocket Secure (WSS) connections, secured with quantum-resistant TLS 1.3 encryption. It was architected to maximize throughput per unit of CPU and network bandwidth, minimize per-transaction overhead, and maintain reliability guarantees across unreliable network conditions—all without depending on any external messaging infrastructure. The result is a distribution layer that is faster, leaner, and more secure than anything achievable through generic JDBC connections or third-party brokers.

The quantum-resistant encryption element deserves specific mention. As the industry prepares for the eventual arrival of cryptographically relevant quantum computing, the data moving through replication pipelines today—financial transactions, healthcare records, sensitive operational data—must be protected against forward-secrecy attacks. GoldenGate’s adoption of post-quantum cryptographic standards in its transport layer is a future-proofing investment that no competing platform currently matches.

Oracle Kernel Integration: A Structural Advantage

No other data integration tool integrates into the Oracle Database kernel the way GoldenGate does. This is not a matter of degree—it is a structural impossibility for any third party. Oracle exposes private, high-speed kernel APIs to GoldenGate that are not available through any public interface (including XStream or LogMiner). The SGA internals and GoldenGate integrated capture APIs have been co-developed across every Oracle Database release since 11.2.0.4 and represent a depth of integration that has been fifteen-plus years in the making and that no competitor can access.



Best → GoldenGate Native

- deepest integration, most parallel and most secure

Great → XStream Out APIs

- good performance, guaranteed ordering
- OEM'd by partners like: Snowflake, IBM/Confluent etc.

Free → Logminer APIs

- single-threaded, order/exact-once not guaranteed by Oracle

Dangerous → Binary Log Readers


- poor security, possible corruption, formats may change at any time

In practice, this means that when you run GoldenGate against an Oracle Database, you are getting a fundamentally different—and fundamentally superior—capture experience compared to any alternative. GoldenGate reads directly from the database kernel's internal transaction buffers at speeds that no JDBC-based tool can approach. The parallel apply engine on the replicat side writes into Oracle Database targets with a level of parallelism and dependency awareness that is only possible because GoldenGate and Oracle Database are, in the deepest sense, the same product. IBM IIDER has solid Oracle support via a standing partnership around Oracle XStream Out. Debezium has functional Oracle LogMiner-based capture and also support for XStream that will require end-users to license GoldenGate anyway. Qlik Replicate and AWS DMS cover Oracle competently but raise significant security concerns over the exposure of Oracle Master Encryption Keys (MEK) to the people working on CDC and data integration.

Ultimately, none of these tools are inside the kernel. GoldenGate is.

Fully Managed Multicloud: Not a Marketplace VM

The final competitive dimension is perhaps the one most frequently misunderstood. Every major cloud provider has a marketplace where enterprise software vendors can offer their products as virtual machine images for self-managed deployment. This is a reasonable way to make software available in the cloud, and most of GoldenGate's competitors offer their products in exactly this way—as VMs you provision, configure, maintain, patch, monitor, and scale yourself.



OCI GoldenGate is something categorically different. It is a fully native managed service, engineered and operated by Oracle alongside the cloud service providers—not a VM in a marketplace. The service runs on Exascale compute infrastructure, co-located within the physical data centers of Azure, Google Cloud, and OCI's own global regions. It is operated by Oracle's cloud DevOps organization with

the same rigor applied to Oracle's own SaaS products. Patching, upgrades, scaling, failover, and monitoring are handled by Oracle—not by the customer's DBA team.

What this means for an architect making a deployment decision is concrete and consequential. Choosing OCI GoldenGate over a marketplace-deployed alternative is the difference between subscribing to a service and running a software product. The infrastructure, the operational expertise, and the SLA are Oracle's problem, not yours. For organizations whose core competency is not running replication infrastructure—which is to say, for virtually every organization—this distinction is the difference between a platform that requires dedicated engineering resources to operate and one that simply works.

The question is not whether GoldenGate's competitors are capable tools. Many of them are. The question is whether they are built for the scale, the criticality, and the AI ambitions of the enterprise you are building—and whether they will still be the right answer in five years.

30 Years of Innovation

When GoldenGate Software was founded in 1995, there were only a small handful of IT systems so important that they could never fail. Banking networks, airline reservation systems, and a few military applications represented the outer edge of continuous availability requirements. Today, that expectation applies to nearly every enterprise system—especially those powering AI-driven applications, digital commerce, and real-time customer experiences.

The original foundation for innovation at GoldenGate Software was the concept of “transaction-safe, logical replication”—simply, the ability to replay database-quality transactions across wide-area networks with zero data loss and guaranteed ordering. That founding principle remains the bedrock of GoldenGate's architecture three decades later, even as the platform has evolved far beyond its origins.

Continuous Availability and Disaster Recovery

The emergence of logical replication was transformational because it allowed IT teams to separate data stores far enough apart to operate on different power grids, different data centers, even different continents—while maintaining transactional consistency at all times. Initially the Tandem NonStop database was central to the startup's growth, but over time GoldenGate expanded to include Oracle Database, SQL Server, IBM DB2/z, DB2/i, DB2/LUW, MySQL, PostgreSQL, and dozens of others.

This heritage of zero-downtime continuous availability remains a core strength of GoldenGate today. Most consumers have personally benefited from GoldenGate high availability solutions at their banks, airlines, and grocery stores—without ever knowing it. GoldenGate is the only data integration platform certified for Platinum and Diamond tier High Availability for Oracle Exadata systems, and it is the de facto enterprise standard for mission-critical replication worldwide.




— Active-Active / Bi-Directional —

- **Highest Availability:** recovery time objective (RTO) of 0-10 seconds, near-zero data loss
- **Active-Active:** all sites active at the same time, built-in conflict detection and resolution and support for blue-green deployment patterns
- **All Oracle Databases:** any version, any deployment, any infrastructure, no shard keys
- **Real-world Performance:**
 - >5TB of REDO/hr at large European bank
 - >1M TPS at global tech manufacturer
 - Sub-second replication at top 5 retailer
 - 16-way bi-directional at Canadian FSI

Trickle-Feed Analytic Data Warehouses

Take yourself back in time to the late 2000s and you will recall the rise of the data warehouse appliance. Vendors such as Teradata, Netezza, Oracle, IBM, Microsoft, and others were competing fiercely to sell pre-engineered systems optimized for analytics. This created a new challenge for IT: how to feed data continuously into these new systems without affecting the performance of the underlying transaction applications.



Traditional ETL tools were fine for many use cases, but for the most important applications there simply wasn't enough capacity to support intensive bulk-extract jobs or enough time in the batch

processing window. Data replication to the rescue. Because GoldenGate extracts data from change logs rather than SQL queries or bulk APIs, it is minimally invasive on the source application and preserves the performance that end-users depend on.

The second great expansion of GoldenGate's market presence came precisely from this capability—providing customers worldwide with continuous, low-latency data feeds to their analytics estates.

Streaming Pipelines and the AI Era

The third and most profound phase of GoldenGate's evolution is unfolding now. The rise of cloud computing, multicloud architectures, open data lakehouses, and most importantly artificial intelligence has created a new imperative: data must be fresh, trusted, and available everywhere in real-time—or AI models will be working with yesterday's reality, not today's. The GoldenGate platform has pivoted decisively to serve as the real-time data fabric layer that makes AI transformation possible.

AI is only as powerful as your data is fresh and trustworthy.

GoldenGate 26ai directly addresses all of these challenges, positioning the platform not merely as a replication tool but as the enterprise's mission-critical AI data fabric (a real-time layer connecting operational systems to analytics and AI).

Everyday Moments by GoldenGate

By the time you read this paper, chances are that you have already generated several GoldenGate transactions in your daily life. Here are just a few examples of everyday activities that create GoldenGate events behind the scenes:

- Buy a latte from your favorite Seattle-based coffee chain
- Update your profile on your professional networking site
- Stop by the grocery store for milk or a quick snack
- Download the latest song stuck in your head
- Have your paycheck processed via Paychex or ADP
- Pay for a recent purchase with PayPal or Apple Pay
- Book a reservation on any major airline
- Withdraw cash from just about any ATM
- Stream your favorite show or browse personalized recommendations
- Submit a cloud-based AI query powered by a real-time vector index

You and your family are already benefiting from GoldenGate daily. And as AI-driven applications become the norm rather than the exception, the footprint of real-time data movement will grow even larger. What will the next 30 years bring?


Proof Points: The Numbers Behind the Story

Claims about enterprise software are easy to make. What separates GoldenGate from the field is a thirty-year record of production deployments at a scale and criticality that most platforms never encounter. The following figures are not projections or benchmarks run in controlled labs—they are the actual operational realities of GoldenGate customers running live systems today.

Scale & Market Reach

- Approximately 80% of the Fortune 500 run GoldenGate in production [5]
- More than 8,000 customers in 150+ countries; across banking, telecommunications, retail, manufacturing, government, healthcare and high technology industries
- More than 20,000 fully-managed OCI GoldenGate deployments; available by subscription in 150+ cloud regions all over the world
- More than 2,000 customers operating GoldenGate in cloud, multicloud, hybrid, and edge deployments
- Over 1,000 certified combinations of databases, operating systems, chipsets, and cloud environments
- GoldenGate supports '*never down*' payment gateways that underpin the vast majority of global commerce 365x7x24hrs — sub-second replication across continents and oceans

Real-World Throughput



Numbers matter in data infrastructure, and GoldenGate's production customers have set benchmarks that redefine what enterprise replication means at the high end:

- A tier-one global bank captures more than 4 TB of Oracle Database transaction logs per hour in a continuous GoldenGate pipeline—without measurable impact on transaction system performance

- A global manufacturing enterprise replicates over 1 million transactions per second across distributed GoldenGate deployments spanning multiple continents
- Oracle tests with GoldenGate 26ai on a quarter-rack Exadata show more than 2.5 million rows per second of changed data processed with Parallel Replicat
- Top social media sites (multiple, actually) use GoldenGate to replicate petabytes of data per day from Oracle transaction logs into Apache Kafka at real-time latencies—on some of the largest known CDC deployments in the industry
- In typical deployments, end-to-end latency from source commit to target apply is measured in milliseconds to single-digit seconds, not minutes or hours

Customer Reference Points

A whitepaper cannot substitute for a conversation with a reference customer, but the breadth of GoldenGate’s deployment base is itself a proof point. The platform runs high availability and real-time data integration for institutions whose names you recognize without further introduction: global money-center banks that clear trillions of dollars daily, airlines whose booking systems handle millions of seat transactions per hour, and technology companies whose platforms serve billions of users. Public references include American Airlines and Starbucks among dozens of others. The platform also serves as the internal backbone for Oracle’s own SaaS applications—arguably the most demanding test of all for a data integration platform.

These numbers tell a consistent story: GoldenGate is not merely a tool that handles demanding workloads—it is the tool that the world’s most demanding organizations have standardized on precisely because alternatives could not meet their requirements. That track record is the foundation on which every new capability described in this paper is built. The next two sections provide a practical decision guide and a set of common architecture patterns before describing those capabilities in detail.

When to Use GoldenGate


GoldenGate is purpose-built for high-stakes, continuous data movement. The table below is a practical guide: if your workload appears in the “Yes” rows, GoldenGate is likely the correct answer. If it appears in the “No” row, a simpler and less expensive tool is probably the better fit.

Scenario	GoldenGate?	Why
Mission-critical high availability and disaster recovery	Yes	Exactly-once delivery, certified MAA Platinum and Diamond tiers—no competing tool qualifies
Real-time AI pipelines (RAG, vector search, agent grounding)	Yes	Sub-second latency, inline embedding generation, native Oracle AI Vector Search integration

Scenario	GoldenGate?	Why
Zero-downtime database migration	Yes	Continuous sync with online cutover; no maintenance window required
Real-time lakehouse or analytics feeds (Iceberg, Snowflake, Databricks)	Yes	ZeroETL Mirror with automatic schema propagation; no batch window
CDC source for Kafka / event-driven architecture	Yes	GoldenGate Data Streams (AsyncAPI) replaces Kafka for many use cases; also feeds Kafka where it exists
High-volume Oracle Database workloads (Oracle-to-Oracle or Oracle-to-cloud)	Yes	Kernel-level integration unavailable to any alternative; per-PDB capture; Autonomous Database support
Simple batch ETL with no latency requirement	No	Purpose-built batch tools are simpler and cheaper; GoldenGate's strengths are unnecessary here
Low-volume development or test database sync	Maybe	GoldenGate Free removes the cost barrier; evaluate based on operational complexity tolerance


Common GoldenGate Architecture Patterns

GoldenGate supports many deployment topologies and use cases. The graphic below highlights the “Top 7” use case patterns. In this section we explain five of them in more detail.




These patterns are not mutually exclusive, a single GoldenGate deployment frequently serves all of them simultaneously. In larger enterprise customers it is very common to see GoldenGate deployed across a broad ecosystem, helping to solve data engineering challenges in each of these use cases.

Pattern 1: Operational High Availability



Active-active and bi-directional. Primary DB “A” ↔ GoldenGate ↔ Primary DB “B”. GoldenGate continuously mirrors every committed transaction from the primary database to one or more additional primaries—across data centers, cloud regions, or continents—with sub-second lag and exactly-once delivery. Failover can be completed in milliseconds and any transaction collisions can be handled seamlessly. This is GoldenGate’s original use case with Tandem NonStop and remains the standard for Oracle MAA Platinum and Diamond-certified deployments.


Pattern 2: CDC to any Kafka-compatible Topics



Source System → GoldenGate Extract → GoldenGate Replicat → Any Kafka-compatible topics. GoldenGate publishes committed transactions as ordered, exactly-once events for downstream consumers that are consuming from Kafka-type topics (or other messaging systems such as Pulsar, RabbitMQ etc). Where Kafka already exists, or when customers need a “fan-out” style distribution layer, GoldenGate serves as a high-fidelity CDC producer and Kafka as the distributor. This pattern is the foundation for event-driven microservices, real-time dashboards, and operational data products.


As an interesting point of fact, GoldenGate originally teamed with LinkedIn engineers to create this CDC-to-Kafka pattern before Kafka was even an Apache standard. Today, GoldenGate produces petabytes of data into some of the world’s largest Kafka clusters across the globe.

Pattern 3: Native, Pub-Sub Data Streams



Source DB → GoldenGate Extract → AsyncAPI Pub-Sub → Any Consumer. As an alternative to Kafka, GoldenGate Data Streams are lower latency, higher volume and more transaction-safe than the general pattern of putting CDC events into Kafka topics. The OpenAPI standards-based implementation of AsyncAPI provides a JSON-friendly REST end-point for any human or AI agent to easily connect and consume a stream of data events from all supported sources. This is the pattern used by organizations that want their consumers to achieve better performance, more trusted data or simpler connectivity than Apache Kafka-based distribution of CDC events.


Pattern 4: Real-Time Lakehouse Ingestion



Source DB → GoldenGate Extract → ZeroETL Mirror → Iceberg / Snowflake / Databricks / Oracle AI Lakehouse. GoldenGate streams row-level changes into open table formats continuously, eliminating the nightly ETL batch window. Schema changes at the source propagate automatically. Analytics and AI models always work against data that is seconds old, not hours. This is the pattern used by organizations replacing legacy ETL infrastructure with a real-time lakehouse architecture.

GoldenGate features unique options for streaming and micro-batch ingestion into Lakehouse systems. In fact, GoldenGate can also aggregate high-volume transactions to reduce compute costs on the target.

Pattern 5: Real-Time AI Pipeline



Source → GoldenGate Extract → GoldenGate AI Microservice → Embedding Generation and other AI enrichments → any GoldenGate Target. GoldenGate intercepts the transaction stream, generates vector embeddings for changed records via OCI Generative AI or any compatible LLM, and writes them to a vector store in real time. The result is an AI knowledge base that reflects operational data within seconds of each committed change—the prerequisite for grounded, accurate agentic AI. This pattern is described in full in the reference architecture section of this paper.

Innovation for the Next 30 Years


The world today is a fundamentally different place than it was in 1995, or even 2020. Almost every industry runs on software, nearly every enterprise is dependent on cloud, and now artificial intelligence is reshaping how organizations make decisions, engage customers, and run operations. Networks and databases are still the lifeblood of software, but we are now so deeply interconnected across multicloud environments, real-time APIs, and AI-driven agents that the demands placed on data infrastructure have never been greater.

It is this intensifying dependence on real-time, trusted data—combined with the explosive growth of AI workloads—that is driving the most ambitious wave of innovation in GoldenGate’s history. The following sections describe the key areas of current and existing innovations that will define the platform for decades to come.

1. Data Fabric & the Backbone for Enterprise AI Data

There has been a genuine explosion of data since the original GoldenGate whitepaper was written. The number of data stores, cloud services, SaaS applications, AI models, and streaming systems that enterprises must integrate has grown by orders of magnitude—and continues to accelerate. In the short time since 2017, GoldenGate has added support for more new platforms than in all of the preceding twenty years combined. But the innovation story no longer ends with connectivity. Today, GoldenGate is the foundation for an enterprise data fabric.

A data fabric is a unified, managed layer (on-premises, cloud, or hybrid) that connects data producers to data consumers in real-time across any infrastructure. GoldenGate serves as the control plane for this fabric—capturing change data from any source, transforming and enriching it in motion, and delivering it to any destination with exactly-once guarantees and full transactional integrity. The result is an organization where every system—from an on-premises Oracle Database to a cloud-native AI lakehouse—is seamlessly connected and continuously synchronized.



This capability has expanded far beyond traditional relational databases. GoldenGate now supports:

- Oracle Autonomous AI Lakehouse and Oracle AI Data Platform
- Apache Iceberg tables on all major cloud platforms
- Databricks Delta Lake and Snowpipe Streaming
- Amazon Kinesis, Redshift, S3, S3 Iceberg Tables and all RDS databases
- Azure SQL, Azure Event Stream, and Microsoft Fabric
- Google Cloud Spanner and BigQuery
- Apache Kafka, Kafka Streams, and OCI Streaming
- Vector databases, NoSQL stores, and AI-native data platforms
- Oracle Applications including E-Business Suite, Retail Cloud, and Transportation Management

The GoldenGate data fabric does not just move data—it creates real-time data products. These are curated, governed, and continuously updated datasets that can be consumed by AI models, analytics platforms, and operational applications alike.

Leading organizations are standardizing on GoldenGate as their enterprise-wide solution for exactly this reason: **one platform, any source, any target, any cloud.**

2. Innovation with Oracle AI Database

As described earlier in this paper, Oracle GoldenGate integrates directly with the Oracle Database kernel through co-developed internal APIs that are not available to third-party tools. This level of integration, refined over more than fifteen years of joint development, underpins GoldenGate's ability to deliver high-fidelity, transactionally consistent data movement for mission-critical systems. GoldenGate is certified for Oracle Maximum Availability Architecture (MAA) at both the Platinum and Diamond tiers—the highest levels of Oracle's availability framework. These tiers require deep integration with Oracle's internal recovery, consistency, and failover mechanisms. As a result, GoldenGate is uniquely positioned among replication technologies to support the most demanding operational workloads, including financial settlement systems, global payment processing platforms, and tier-one transactional databases.

GoldenGate also aligns with the security posture of Oracle AI Database. It supports AES-256 encryption for data at rest and leverages post-quantum–resilient cryptographic mechanisms, including ML-KEM within TLS 1.3, for data in motion [7]. As organizations prepare for cryptographically relevant quantum computing, securing the replication pipeline—often one of the largest and least protected data movement surfaces—becomes a critical requirement.

Within Oracle Multitenant environments, GoldenGate provides precise, per-Pluggable Database (PDB) capture. This allows change data to be scoped to individual tenants rather than mined at the Container Database (CDB) level, preserving isolation boundaries and simplifying compliance in multi-tenant deployments. GoldenGate is also the only replication platform currently supported for change data capture from Oracle Autonomous Database, where system-level access is restricted and alternative approaches are not viable.

GoldenGate extends its integration to newer Oracle data models as well. It captures changes from JSON Relational Duality Views directly from the transaction log and publishes them as complete, ACID-consistent business objects via GoldenGate Data Streams. Rather than emitting only row-level deltas, GoldenGate can propagate fully materialized document representations, enabling downstream systems to consume data in the same form used by applications.

Starting with Oracle AI Database 26ai, GoldenGate also supports replication of blockchain and immutable tables, including the associated metadata required for auditability and verification. These append-only, tamper-evident structures are increasingly used in compliance-sensitive environments, and correct replication requires awareness of their specialized semantics.

For Oracle AI Vector Search, GoldenGate provides native support for real-time vector data pipelines. It captures and replicates the Oracle VECTOR data type, supports inline embedding generation during replication via integration with OCI Generative AI, and enables construction of real-time vector hubs across active-active and multicloud architectures. Using the DBFUNCTION feature, Replicat processes can invoke embedding models directly within the database, eliminating the need for separate transformation pipelines. In addition, Oracle Machine Learning and ONNX inference can be applied within GoldenGate Stream Analytics, allowing AI-driven enrichment and decision-making to occur directly within the streaming layer.



GoldenGate is also integrated with Oracle AI Database Private Agent Factory and Private AI Container, enabling AI agents to operate on continuously synchronized, transactionally consistent operational data. GoldenGate's exactly-once delivery model is particularly important in this context, as it ensures that agents do not act on duplicated or out-of-order events.

Finally, GoldenGate's downstream mining capability provides a practical path for modernizing legacy environments. Organizations running older Oracle Database versions can forward redo logs to a newer Oracle AI Database instance, where GoldenGate performs mining and replication. This allows replication infrastructure and downstream systems to operate on current database releases without requiring changes to the source application. In effect, the data pipeline can be modernized independently of the application lifecycle—an approach made possible by the deep integration between GoldenGate and Oracle Database.

3. Multicloud Deployment Everywhere

GoldenGate did not achieve market dominance by being exclusive to Oracle. The platform's DNA—rooted in the Tandem NonStop database from the very beginning—has always embraced heterogeneity. Today, that philosophy extends to the most ambitious multicloud program in the enterprise software industry.



Oracle Cloud Infrastructure now operates more than 200 live and planned regions globally. GoldenGate runs natively across all of them—as a fully managed OCI GoldenGate service, as self-managed software in the OCI Marketplace, or as a containerized deployment in any Kubernetes environment. More importantly, the platform extends seamlessly beyond Oracle's own cloud [8]:

- Oracle Database@Azure: >33 regions with OCI GoldenGate as a fully managed service
- Oracle Database@Google Cloud: >5 regions with fully-managed GoldenGate services
- Oracle Database@AWS: long standing support for self-managed deployments on EC2

- On-premises: Linux, Windows, Mainframe (IBM z/OS), and other hardware platforms
- Edge Devices: certified with Oracle Roving Edge devices for disconnected or IoT scenarios

The operational simplicity of this multicloud story is as important as the breadth. For decades, managing GoldenGate across large numbers of database hosts was a legitimate administrative challenge—particularly for organizations running DB2 for z/OS mainframe systems where MIPS cost savings from off-host deployment were significant. The modern OCI GoldenGate service eliminates that complexity entirely, providing a unified management plane regardless of where the data lives. Whether your source database is in Tokyo or your AI model training is in Frankfurt, GoldenGate routes the transactions with sub-second latency and transactional integrity intact.

4. Autonomous & AI-Assisted Operations

For the vast majority of GoldenGate's first two decades, the engineering focus was on making the platform rock-solid for disaster recovery and high availability. The command-line interfaces and parameter files of that era were entirely appropriate for the DBAs and developers who managed those systems—and GoldenGate thrived, even as competitors offered prettier surfaces with worse performance.

Starting with GoldenGate 12.3, simplification became a first-class engineering priority. Browser-based administration, REST APIs, automated conflict detection and resolution, and cloud-native deployment options all arrived in rapid succession. Today, with GoldenGate 26ai, that simplification journey has reached a new frontier: the platform is actively using artificial intelligence to make itself easier to operate.

The GoldenGate AI Microservice—now a standard option in GoldenGate 26ai deployments—can work with local or remote large language models to bring generative AI capabilities directly into the replication lifecycle. This opens entirely new possibilities:

- Natural language interfaces for creating and managing GoldenGate pipelines
- AI-assisted monitoring and triage: the platform explains anomalies in plain English
- Intelligent tuning recommendations based on workload patterns
- GoldenGate operations engineers will just speak to the service (via MCP tools and skills) to tell it what to do and how to behave
- Auto-schema management: target tables automatically kept in sync with source changes
- LLM and GPT operator integration within Stream Analytics pipelines for real-time data enrichment

We have now entered an era where data replication customers can get Ferrari-level performance and Toyota-level reliability, with a cloud-native subscription service that makes the whole platform as accessible as any modern SaaS application. Oracle GoldenGate makes this possible right now—and the AI-assisted operations layer is only the beginning.



5. Data Mesh Service Topologies

At the heart of GoldenGate's modern DNA is a microservices architecture that was reconceived from the ground up starting in 2016—before 'microservices' had even entered the mainstream IT lexicon. What began as a Service Component Architecture has evolved into a fully realized, HTTP/REST-native platform where Administration, Distribution, Metrics, and now AI services each operate as discrete, independently scalable building blocks – aka: a Service Mesh for real-time data.

Unlike other data integration tools, GoldenGate utilizes a true mesh topology. A sequence of data events captured from one particular source system can be transmitted along many separate 'hops' through different GoldenGate nodes – much like a WiFi mesh utilizes nodes as repeaters for the signal. These architectural foundations (*service mesh deployments and network mesh routing*) make GoldenGate uniquely suited to serve as the strategic control plane for enterprise data movement—not just a point-to-point replication tool. Consider what this means in practice: a single GoldenGate deployment can simultaneously capture transactions from an on-premises Oracle Database, route them through a global data stream, enrich them with AI-generated embeddings, apply them to an Iceberg lakehouse in Azure, and publish them as a JSON data stream to downstream analytics tools—all with exactly-once guarantees and full governance.

GoldenGate Data Streams, introduced with the 23ai release and enhanced in 26ai, provides a Pub/Sub interface based on AsyncAPI that any application or service can subscribe to in real-time. Combined with Oracle's JSON Relational Duality Views, this creates a remarkable new pattern: producer applications work with business objects in JSON format, the data is stored efficiently in relational row format, and any change automatically propagates to subscribers as a complete, consistent document. This eliminates the need for separate Kafka infrastructure for many event-driven application architectures.

What if we could combine database events and streaming events from every system across a company—on-premises, cloud, and multicloud—into a unified, real-time data fabric that powers both operational data (runs the business) and AI-driven analytical decision making simultaneously? That is exactly what GoldenGate delivers today.

6. Containers, Kubernetes & Developer-First Deployment

One of the most profound technology shifts to affect enterprise software in recent years has been the broad adoption of container-based deployment. Docker containers allow complex application software to be pre-configured, packaged, and shipped to any host platform as a single deployable unit. Kubernetes orchestrates those containers at enterprise scale, providing automated scheduling, self-healing, and horizontal scaling across any cloud or on-premises infrastructure.

GoldenGate embraced this shift early and has continued to deepen its container-native capabilities with every major release. With GoldenGate 26ai, Oracle has introduced full container downloads that enable customers to run "GoldenGate as a Service" in their own infrastructure—whether that's a private Docker environment, a Kubernetes cluster, or an OpenShift deployment. This developer-first




experience dramatically reduces the time to get started, lowers operational overhead, and makes it practical to scale GoldenGate horizontally to match workload demands.

The practical impact of this approach is significant for organizations across every industry. A financial services firm can now spin up a new GoldenGate deployment in minutes—containerized, version-controlled, and fully integrated with their Kubernetes orchestration layer—rather than the multi-day process of traditional software installation. A retail enterprise can run identical GoldenGate containers across three cloud providers simultaneously, ensuring that their real-time inventory data is consistent everywhere it is needed.

For developers building on the GoldenGate platform, the REST-native microservices architecture means that every GoldenGate capability can be controlled programmatically over HTTP—enabling integration with CI/CD pipelines, infrastructure-as-code tools, and any orchestration framework. This is the “runs everywhere” promise made operational.

7. Security, Observability, Governance & Trust



GoldenGate began its journey in the financial services industry, for banking and ATM cash networks. In these environments data is literally money, and the software tasked with processing these financial transactions has had to be among the very most secure software in the world. This is the context in which GoldenGate has thrived, continuously innovating and implementing one of the most secure data integration ecosystems for sensitive data. In 2026, GoldenGate introduced a series of quantum-resistant algorithms for protecting data-in-motion as well as data-at-rest.


In any distributed software deployment, comprehensive monitoring is both essential and challenging. GoldenGate’s classic architecture—an agent-based system that may have hundreds or even thousands of installations across different host servers worldwide—makes real-time observability across such large, distributed deployments a significant engineering problem. The microservices

architecture has fundamentally changed this equation.

The GoldenGate Metrics Service provides real-time monitoring via RESTful APIs with a lightweight local database storing up to one year of metrics per service instance. This open framework allows any GoldenGate customer to integrate with the monitoring and observability tools of their choice—from the ELK stack to Grafana, Prometheus, Oracle Enterprise Manager, and beyond. Enterprise teams that have standardized on OCI’s integrated observability tooling get a seamlessly unified view of their entire GoldenGate estate alongside all other OCI services.

But in the AI era, observability must mean more than uptime and lag metrics. Data governance and auditability have become strategic requirements. Enterprises using GoldenGate to power AI training pipelines, regulatory reporting systems, or financial transaction processing need to be able to answer the question: “Is this data exactly what it should be?”

GoldenGate Veridata—now available both on-premises and as part of OCI GoldenGate following the 26ai release—provides algorithmic verification that data is exactly the same across distributed systems, providing a mathematically guaranteed audit trail for the most demanding compliance scenarios.




The platform also provides air-gap replication capabilities for high-security networks—ensuring that even in the most isolated, regulated IT environments, GoldenGate can move data with maximum security for data in motion and at rest. This combination of real-time observability, algorithmic data verification, and high-security deployment options makes GoldenGate the most trustworthy data movement platform available.



8. GoldenGate for AI & Agentic Workloads

During most of the first twenty-five years of data replication, the focus was narrowly defined on the capture and movement of data events. GoldenGate technology became so optimized at this core function that transactions could be seen anywhere in the world with very low latency or lag. What if we were to combine database events from all enterprise data stores in a single stream, enrich them with AI-generated intelligence, and use them to power both analytics and agentic AI systems simultaneously?



That is the ambition of GoldenGate 26ai's AI-native capabilities. The GoldenGate AI Microservice, introduced in this release, supports integration with OCI AI Generative AI, OpenAI (GPT-4 Turbo, GPT-3.5-Turbo, and the full embeddings family), Azure OpenAI, and any custom API-compatible LLM. This means that GoldenGate can now:

- Capture real-time transactions from any database source
- Generate vector embeddings for changed data in motion
- Enrich data events with AI-generated metadata, classifications, or summaries
- Write enriched embeddings directly to vector databases for RAG and GraphRAG pipelines
- Expose GoldenGate pipelines as tools for OCI Generative AI agents
- Integrate with Oracle Integration Cloud AI Agent Framework for reactive, event-driven AI workflows

Stream Analytics 26ai extends this AI capability into the stream processing layer, adding AutoML integration with Oracle AI Database 26ai, LLM and GPT operators for data quality and enrichment within streaming pipelines, and vector embedding generation as part of the streaming workflow. The

result is a platform where AI is not bolted on after the fact—it is woven into the fabric of data movement itself.

The strategic value of this cannot be overstated. Real-time intelligence eliminates the data deserts that plague AI deployments today—where models trained on stale data make decisions that reflect a reality from hours or days ago. Trusted, real-time data delivered by GoldenGate allows AI to sense and respond to real events instantly, enabling the automated decision making that defines the next generation of enterprise AI.

9. Simple, Fully Managed Multicloud Platform

Perhaps the most transformational shift at Oracle over the past decade has been the company's commitment to cloud-native, autonomous services. GoldenGate has been an instrumental part of this transformation. Behind the scenes, GoldenGate provides much of the resiliency of Oracle's own SaaS applications and cloud platform. For customers, OCI GoldenGate is the most capable managed data integration service available on any cloud.

OCI GoldenGate is available in all Oracle Cloud regions as a fully managed service, delivering a 99.95% SLA [5] with automatic scaling, patching, and optimization managed entirely by Oracle. The autonomous nature of the service means customers can focus on the business outcomes of their data pipelines rather than the operational mechanics of running the software. For organizations that have standardized on Oracle Cloud, the integration with OCI IAM, OCI Streaming, Oracle Autonomous Database, and Oracle AI Data Platform is seamless and deeply optimized.

The power of this approach mirrors what Microsoft customers experience with Office 365—a simple cloud subscription grants the ability to use a world-class enterprise platform anywhere, without managing infrastructure. Importantly, the endpoints of the data replication can be in the most optimal networks for each use case: it is not required to route data through Oracle's networks at all when connecting, say, an on-premises data center to an Amazon cloud target. GoldenGate's architecture ensures that performance, latency, and security decisions remain with the customer.

10. Inexpensive, Consumption-Based & Free Pricing

Probably the most persistent and misunderstood rumor about GoldenGate over the years has been that it is expensive. The reality has always been more nuanced. GoldenGate's established pricing has been very fair for mission-critical use cases where the cost of the software is fractional compared to the overall cost of making an application available 99.999% of the time. But as the popularity of real-time data integration has grown for analytics, AI, and data migration use cases, the need for flexible, low-cost consumption-based pricing has grown correspondingly.

Oracle has responded with two important initiatives. First, GoldenGate Free provides a no-cost entry point for developers, data engineers, and organizations that want to explore the platform's capabilities without a license commitment. This is a meaningful change for the GoldenGate

ecosystem: it lowers the barrier to adoption, accelerates learning, and creates a new generation of GoldenGate-fluent practitioners.

ORACLE

Oracle Container Registry

Sign In

GoldenGate Repositories

Short URL : <https://container-registry.oracle.com/ords/ocr/ba/goldengate>

Repository	Description	Open Source License Text
goldengate-free	Oracle GoldenGate Free	The container image you have selected and all of the software that it contains is licensed under Oracle Free Use Terms and Conditions that are provided in the container image. Your use of the container is subject to the terms of those licenses.
goldengate-mysql-free	Oracle GoldenGate for MySQL Free	The container image you have selected and all of the software that it contains is licensed under Oracle Free Use Terms and Conditions that are provided in the container image. Your use of the container is subject to the terms of those licenses.
goldengate-oracle-free	Oracle GoldenGate for Oracle Free	The container image you have selected and all of the software that it contains is licensed under Oracle Free Use Terms and Conditions that are provided in the container image. Your use of the container is subject to the terms of those licenses.
goldengate-postgresql-free	Oracle GoldenGate for PostgreSQL Free	The container image you have selected and all of the software that it contains is licensed under Oracle Free Use Terms and Conditions that are provided in the container image. Your use of the container is subject to the terms of those licenses.
goldengate-sqlserver-free	Oracle GoldenGate for SQL Server Free	The container image you have selected and all of the software that it contains is licensed under Oracle Free Use Terms and Conditions that are provided in the container image. Your use of the container is subject to the terms of those licenses.
goldengate-studio-free	Oracle GoldenGate for Studio Free	The container image you have selected and all of the software that it contains is licensed under Oracle Free Use Terms and Conditions that are provided in the container image. Your use of the container is subject to the terms of those licenses.

1 - 6

Second, OCI GoldenGate’s elastic, consumption-based pricing model aligns cost directly with usage—making it practical and cost-effective for workloads that are temporary, highly variable, or just getting started. In many scenarios, including data migrations, proof-of-concept deployments, and AI pipeline development, this model provides a dramatically better Total Cost of Ownership (TCO) compared to traditional perpetual licensing. The GoldenGate platform has become truly accessible to organizations of every size, at every stage of their data journey.

Reference Architecture: A Real-World Pipeline

Innovation sections are useful for understanding what a platform can do. A reference architecture is useful for understanding how it actually does it. The pipeline described below is representative of deployments that GoldenGate runs in production at tier-one financial institutions today. It is not a simplified toy example—it is the kind of architecture that a senior data engineer or solutions architect would recognize immediately as production-grade.


The Scenario: Real-Time Financial Intelligence

Imagine a tier-one bank with a core banking application on Oracle Exadata and payments processing running on NonStop Enscribe (Base24). Every second, tens of thousands of transactions are committed—deposits, withdrawals, transfers, loan payments, and card authorizations. At peak times, there could be more than a million rows of data per second being processed.



The bank's strategic objectives require four things that a traditional architecture cannot provide simultaneously: data stores that can never go down (even if an entire power grid is lost), a real-time fraud detection model that operates on live transaction patterns, a continuously updated analytical lakehouse for regulatory capital reporting, and a RAG-powered customer service agent that knows the current state of every account. Here is how GoldenGate delivers all four.

The pipeline begins with a GoldenGate Integrated Extract process running on the Oracle Database host. It reads committed transactions directly from the Oracle SGA layer using the private GoldenGate API—the same kernel-level interface that GoldenGate has been optimizing with Oracle for fifteen years. No SQL queries hit the live application tables. No batch jobs compete for I/O. The Extract process reads change records in the order they were committed and transmits them, with full transactional context, via the native binary trail protocol.



The GoldenGate Distribution Service fans the financial transactions out to multiple downstream consumers over quantum-safe TLS 1.3 secure web sockets with ML-KEM for key handling. **Target 1** is the secondary data center for high availability on the source applications. **Target 2** is AWS Kinesis, a Kafka-compatible managed service that receives the full transaction stream for real-time consumers. **Target 3** is a Databricks lakehouse running from Azure, where GoldenGate's ZeroETL Mirror writes Iceberg-formatted tables in a Bronze/Silver/Gold medallion architecture—Bronze receives raw CDC events, Silver applies light transformation and deduplication, Gold provides business-ready analytical tables for the regulatory reporting system. **Target 4** uses the GoldenGate AI Microservice, which intercepts the transaction stream, generates vector embeddings for account state changes using OCI Generative AI, and writes those embeddings to an Oracle AI vector database that powers the CX customer service RAG agent.

Latency Expectations

In a well-tuned GoldenGate deployment of this type, the end-to-end latency from source database commit to availability in all four targets is typically under one second for normal transaction volumes, and under three seconds at peak load. This is not a theoretical bound—it reflects the checkpoint commit interval and network round-trip time between the GoldenGate services, both of which are configurable and predictable.

For the fraud detection use case, this means that a suspicious transaction pattern—say, three card authorizations in different geographies within sixty seconds—can trigger a model evaluation and a hold decision before the fourth authorization arrives. For the regulatory reporting use case, it means that capital adequacy calculations can run on data that is seconds old rather than hours old. For the RAG agent, it means a customer service representative asking “what is the current balance on account X?” receives an answer that reflects the last committed transaction, not last night’s batch extract.

Failure Handling and Recovery

No architecture document is complete without an honest account of failure behavior. GoldenGate’s approach to failure recovery is one of the platform’s most important—and least appreciated—design characteristics.


Every GoldenGate process maintains a checkpoint: a precise, durable record of its position in both the source transaction log and the target apply queue. If a network connection drops mid-stream, GoldenGate does not restart from the beginning of the log or from a time-based approximation. It restarts from the exact checkpoint—picking up precisely where it left off, with no duplicate delivery and no gap in the sequence. This is the architectural basis for exactly-once guarantees, and it works across any type of failure: network partition, process crash, host reboot, or even a rolling database upgrade.

For high availability of the GoldenGate services themselves, the platform supports Oracle’s Maximum Availability Architecture (MAA) patterns. The Distribution Service can be configured in active-active mode (traditional active-passive failover not required), ensuring that a failure of the first primary distribution host results in a seamless and immediate takeover by the second primary instance—typically within milliseconds. For the on-premises GoldenGate MAA deployments, customers can expect 99.999% availability. With the OCI GoldenGate managed deployments, this failover is handled automatically as part of the 99.95% SLA (single instance), with no operator intervention required. Multiple instances and cross-region deployments can be configured to reach 99.99% and 99.999% high availability. The results are real-time data pipelines that are as resilient as the mission-critical databases that GoldenGate serves.

New Frontiers

Decoupling your Data from the Lakehouse

GoldenGate Data Fabric represents the next evolution of Oracle GoldenGate—from a market-leading replication platform into a real-time enterprise data fabric for AI, analytics, and operational workloads. At its core, the platform connects data producers to data consumers across any infrastructure, cloud, or lakehouse architecture with transactional consistency, sub-second latency, and enterprise-grade trust. Rather than treating operational databases, streaming systems, and analytical platforms as isolated silos, GoldenGate Data Fabric creates a unified real-time layer that continuously distributes trusted data products and data events wherever they are needed. This architecture enables organizations to build AI-ready platforms without centralizing all enterprise data into a single monolithic repository.



The platform introduces a modern, AI-native approach to data engineering and real-time integration. GoldenGate Data Fabric combines CDC replication, continuous transformation and loading (CTL), ETL/E-LT pipelines, stream processing, governance, and observability into a single managed environment operating across OCI, Azure, Google Cloud, AWS, on-premises infrastructure, and edge deployments. Users can develop pipelines using natural language prompting, low-code visual tooling, or high-code APIs and DSLs designed for AI-assisted development environments. Unlike traditional data integration products built around batch movement and centralized processing, GoldenGate Data Fabric is purpose-built for event-driven architectures where pipelines execute continuously as business events occur.

A defining architectural principle of GoldenGate Data Fabric is the decoupling of the runtime layer from the underlying storage and lakehouse infrastructure. Organizations can operate across Oracle Database, Apache Spark, Apache Flink, Iceberg lakehouses, and heterogeneous cloud environments without rewriting data pipelines for each platform. This flexibility allows enterprises to optimize

infrastructure cost, avoid long-term vendor lock-in, and deploy workloads wherever latency, governance, sovereignty, or AI processing requirements demand. The platform supports both operational and analytical data products simultaneously, enabling real-time transactional systems, streaming applications, AI vector pipelines, and analytical workloads to coexist within a single governed data fabric architecture.

Most importantly, GoldenGate Data Fabric positions Oracle at the center of the emerging AI data platform market. The platform integrates directly with Oracle AI Database, Oracle Autonomous AI Lakehouse, OCI Generative AI services, and vector-enabled data architectures to create continuously updated, AI-ready data pipelines. Real-time embedding generation, JSON Relational Duality event streaming, governed data products, and open lakehouse support allow enterprises to operationalize AI using fresh, trusted business data rather than stale batch extracts. As organizations move toward agentic AI systems and real-time decision automation, GoldenGate Data Fabric becomes more than a data integration tool—it becomes the operational control plane connecting enterprise transactions, analytical systems, and AI-driven applications into a single real-time intelligence fabric.

An Open and Independent ZeroETL Layer

ZeroETL has emerged as one of the most important architectural patterns in modern data engineering because it removes the operational complexity, latency, and fragility associated with traditional ETL pipelines. But most implementations in the market today come with an important tradeoff: the convenience of ZeroETL is tightly coupled to a single vendor's ecosystem. In practice, this means the source system, runtime layer, storage platform, and analytical services are all expected to live within the same cloud boundary. AWS ZeroETL is optimized for AWS-native targets. Microsoft Fabric Mirroring is designed around Fabric-managed storage and services. These approaches simplify deployment inside a single platform, but they also create long-term architectural gravity that can make multicloud strategy, platform portability, and infrastructure optimization significantly more difficult over time.

A more flexible approach is to separate the operational simplicity of ZeroETL from the underlying lakehouse and cloud infrastructure itself. Instead of forcing organizations to standardize on a single proprietary storage layer or analytics stack, ZeroETL pipelines can operate across heterogeneous environments—streaming operational data continuously into Oracle and non-Oracle databases, Iceberg lakehouses, cloud-native analytical platforms, and distributed AI data architectures simultaneously. This allows organizations to preserve the benefits of real-time automated ingestion without surrendering architectural independence. Pipelines become portable assets rather than cloud-specific implementations, enabling enterprises to optimize for cost, performance, governance, data sovereignty, or AI processing requirements without rewriting the entire integration layer every time infrastructure strategy changes.

This architectural decoupling becomes even more important in the AI era. Enterprises increasingly operate across multiple clouds, multiple analytical engines, and multiple data persistence models simultaneously. Operational data may originate in Oracle Database, stream into Iceberg lakehouses, feed vector pipelines for RAG systems, and support real-time analytics in entirely different



environments—all at the same time. A ZeroETL architecture constrained to a single vendor's storage platform fundamentally limits that flexibility. By contrast, a platform-agnostic approach allows organizations to build a true distributed real-time data fabric where the pipelines, governance layer, and operational semantics remain consistent regardless of where the data ultimately resides. The result is not only lower operational complexity, but also strategic freedom: the ability to evolve infrastructure choices over time without re-engineering the foundation of the data architecture itself.

GoldenGate goes Interplanetary

For most of its thirty-year history, GoldenGate has operated in some of the most demanding and physically distributed environments on Earth. The platform was designed around the idea of decentralized, transaction-aware data movement across unreliable and bandwidth-constrained networks—long before “data mesh” became a mainstream architectural term. Today, GoldenGate already operates successfully in disconnected and semi-disconnected environments including submarines, cruise ships, offshore platforms, remote mining operations, and mobile medical systems such as the Royal Flying Doctor Service supporting remote communities across vast geographic regions. These environments share a common challenge: the network cannot be assumed to be continuously available, latency may be unpredictable, bandwidth may be constrained, and yet the integrity and ordering of operational data still matters. GoldenGate's checkpoint architecture, lightweight binary trail format, and transaction-safe replication model make it exceptionally well suited to exactly these conditions.

Over the next thirty years, humanity will begin extending these same operational requirements beyond Earth itself. Long-duration human spaceflight, lunar habitats, autonomous industrial systems, orbital manufacturing platforms, and eventually permanent installations on Mars will all depend on distributed IT systems operating across extremely high-latency and intermittently connected networks. These environments will require far more than simple file synchronization or eventually consistent object replication. Mission-critical operational systems will still need trusted transactions, ordered event delivery, compact data movement, and resilient synchronization between isolated edge environments and centralized command infrastructure. A future interplanetary data mesh will require software capable of operating intelligently across disconnected, delay-tolerant, and low-bandwidth conditions while preserving transactional integrity whenever connectivity is restored. These are precisely the categories of engineering problems GoldenGate has already spent decades solving on Earth.

In this sense, space represents a natural extension of GoldenGate's architectural trajectory rather than a radical departure from it. The same distributed mesh principles that allow GoldenGate to synchronize operational systems across oceans, remote industrial facilities, and disconnected mobile environments today could ultimately enable trusted data movement between orbital infrastructure, lunar data centers, deep-space vehicles, and planetary habitats tomorrow. As enterprises and governments expand human presence beyond Earth, real-time and transactionally consistent data movement will become foundational infrastructure for space operations, autonomous systems, AI-driven mission control, and interplanetary logistics. The next frontier for GoldenGate may not simply be another cloud region or another continent—it may ultimately be another world.



The Next 30 Years: GoldenGate stays Gold

Take yourself forward in time thirty years. It is nearly impossible to predict with precision what technology will look like—but it is possible to identify the forces that will shape it. The Earth, the Moon and Mars will be more connected, more data-intensive, more AI-dependent and more demanding about the freshness and trustworthiness of information than they are today. Every one of these trends plays directly to GoldenGate's core strengths.

The vision for GoldenGate over the next three decades is built on three bold themes. The first is zero-latency data movement: a future where the concept of a batch window is as obsolete as a mainframe punch card, where every enterprise system is continuously synchronized with sub-millisecond precision. The second is autonomous data management: where GoldenGate not only moves data but understands it, using AI to make intelligent decisions about routing, enrichment, quality, and governance without human intervention. The third is AI-native enterprise infrastructure: where GoldenGate serves as the central nervous system connecting every operational system, every AI model, and every analytics platform into a coherent, real-time intelligence fabric.

Data replication capabilities will be more relevant in the future, not less. The enterprises that win the next three decades will be those that treat real-time, trusted data not as an IT project but as a strategic business imperative.

GoldenGate has spent thirty years earning the trust of the most demanding organizations in the world. That trust—built transaction by transaction, byte by byte—is the most durable competitive advantage in enterprise software. And it is the foundation upon which the next thirty years of innovation will be built.

Limitations & Tradeoffs

A technical paper that presents only strengths is marketing collateral. This one aims to be something more useful: an honest assessment that helps architects and decision-makers understand not just where GoldenGate excels, but where it requires careful planning—and where a different tool might be the right answer. Thirty years of production deployments have clarified these boundaries with a precision that no benchmark can replicate.

The first consideration is operational complexity at scale. For small solutions GoldenGate can be a “set it and forget it” type solution. However, enterprise deployments with dozens of pipelines across multiple source and target systems are sophisticated distributed infrastructure that require ongoing governance: configuration management, trail file housekeeping, checkpoint monitoring and capacity planning for peak transaction loads. Oracle has made significant investments in simplification—the

microservices REST API, the OCI GoldenGate managed service, AI-assisted MCP interfaces in 26ai—and these genuinely reduce the operational burden. But for organizations without experienced GoldenGate practitioners or a willingness to invest in that expertise, the learning curve is real. This is particularly true for large, heterogeneous deployments spanning many databases and cloud environments, where the combinatorial complexity of source/target configurations grows quickly.

The second consideration is cost relative to workload requirements. GoldenGate's enterprise licensing is priced for mission-critical use cases—and for those use cases, where downtime or data loss has direct financial and regulatory consequences, the cost is easily justified. But not every integration scenario carries that level of criticality. A development team synchronizing a test database, a small analytics workload reading infrequently from a low-volume source, or a simple file-based data pipeline with batch-acceptable latency may find that GoldenGate's capabilities—and its cost—exceed what the use case demands. OCI GoldenGate and GoldenGate Free address the lower end of this spectrum meaningfully, but the full enterprise platform is calibrated for enterprise problems.

The third consideration is latency floor and infrastructure overhead. GoldenGate is extraordinarily fast by any reasonable standard—sub-second end-to-end latency in typical configurations is a real and achievable target. But GoldenGate is a distributed system with real network dependencies, checkpoint writes, and trail file I/O. In environments where microsecond-level latency is required—certain high-frequency trading systems or in-memory database mirroring use cases—GoldenGate operates at a different layer of the stack than those specialized tools. It is the right answer for the overwhelming majority of enterprise replication requirements; it is not a substitute for shared-memory or hardware-accelerated synchronization in exotic latency-sensitive architectures.

Every powerful tool has a domain where it is the right answer and a domain where it is not. GoldenGate's domain is broadly defined—but knowing its edges makes you a better architect. None of these tradeoffs diminish GoldenGate's standing as the premier platform for mission-critical, real-time data movement. They are simply the honest contours of a technology that has been refined in production for three decades. Understanding them is what distinguishes a deployment that thrives from one that struggles.

Conclusion

GoldenGate has been one of the most successful enterprise software platforms of the past three decades. It has powered the data backbone of global banking, telecommunications, retail, and technology, and it is now positioned at the center of the most important technology transition since the shift to cloud: the rise of enterprise AI.

Try to imagine a future where it is acceptable for business applications to have more downtime, operate at slower batch driven data cycles, or where it is acceptable to train AI models on stale data—it is not going to happen. The future will bring even more intense demands on application availability, and even more focus on providing AI insights using up-to-the-moment data. Real-time, trusted data movement will be more relevant in the future, not less.

Oracle is investing to ensure that GoldenGate remains the enterprise standard for this mission-critical capability, and that the platform continues to evolve in step with the most demanding requirements of the AI era.

These investments include:

- Data Fabric – for unified, governed, real-time data products across any infrastructure
- Kernel Integration with Oracle AI Database – performance that no competitor can match
- Multicloud Deployment Everywhere – OCI, Azure, Google Cloud, AWS, and edge
- Autonomous & AI-Assisted Operations – simplicity at mission-critical scale
- Real-Time Data Fabric Architecture – the strategic control plane for data movement
- Containers, Kubernetes & Developer-First Deployment – runs anywhere
- Observability, Governance & Trust – auditable, verifiable, exactly-once pipelines
- GoldenGate for AI & Agentic Workloads – embeddings, LLMs, and RAG pipelines
- Fully Managed Multicloud Platform – OCI GoldenGate for fully-managed simplicity
- GoldenGate Free & Elastic Pricing – accessible to everyone at every scale

We are at an extraordinary intersection: thirty years of proven, mission-critical technology meeting the most transformative wave of innovation in the history of enterprise computing. Oracle GoldenGate is at the forefront of this intersection—empowering customers all over the world to build the AI-ready, real-time data platforms that will define the next three decades of competitive advantage.

What to do next—based on where you are in your evaluation:

1. **Exploring:** Start with GoldenGate Free (oracle.com/integration/goldengate/free) — zero cost, no license commitment, full platform capability for development and proof-of-concept workloads.
2. **Evaluating:** Run a guided hands-on lab at Oracle LiveLabs (livelabs.oracle.com) — pre-built environments covering HA replication, lakehouse ingestion, and AI vector pipelines.
3. **Ready to deploy:** Contact Oracle at oracle.com/goldengate to speak with a solutions engineer about production architecture, sizing, and OCI GoldenGate managed service options.



Appendices

Glossary of Terms

Oracle GoldenGate — Oracle's enterprise data integration platform for real-time, transactionally consistent data movement across databases, applications, streams, cloud platforms, lakehouses, and AI systems.

OCI GoldenGate — The fully managed Oracle Cloud Infrastructure service for Oracle GoldenGate. It delivers GoldenGate's capabilities as a managed cloud service, removing the operational burden of provisioning, patching, scaling, and monitoring replication infrastructure.

GoldenGate 26ai — The 2026, AI-era release generation of Oracle GoldenGate, adding native capabilities for AI and agentic workloads, stream processing, AI-assisted operations, and deep integration with Oracle AI Database.

Change Data Capture (CDC) — The process of detecting and moving only changed data from a source system. CDC enables continuous synchronization without repeatedly copying entire datasets, and is the foundation of GoldenGate's real-time data movement.

Log-Based Capture — A CDC technique that reads committed changes from database transaction logs rather than querying application tables. GoldenGate uses log-based capture to minimize source-system impact while preserving full transactional context.

Extract — The GoldenGate process that captures committed changes from a source system. For databases, Extract reads from transaction logs or database-integrated capture mechanisms rather than touching application tables.

Replicat — The GoldenGate process that applies captured changes to a target system, preserving transaction order and applying changes according to the configured delivery and mapping rules.

Trail File — GoldenGate's durable intermediate format for captured change records. Trail files let GoldenGate move, checkpoint, recover, and replay transaction data reliably across distributed environments.

Checkpoint — GoldenGate's durable record of its processing position in the source and target streams. Checkpoints allow recovery from any failure point without losing data or re-delivering already-applied transactions.

Exactly-Once Delivery — The guarantee that each committed transaction is delivered and applied once, with no duplicates and no gaps. GoldenGate provides this property, which is essential for financial systems, operational databases, and other mission-critical workloads.

Commit Order — The sequence in which transactions are committed at the source. GoldenGate preserves commit order end to end, so downstream systems reflect the same business reality, in the same sequence, as the original source.

Transaction Boundary — The point at which a database transaction begins and ends. GoldenGate preserves transaction boundaries so downstream systems receive complete business transactions rather than disconnected row-level fragments.

High Availability (HA) — Architectures designed to keep systems continuously available despite failures. GoldenGate supports HA by synchronizing data across systems, regions, or data centers in real time.

Disaster Recovery (DR) — The ability to restore or continue operations after a major outage. GoldenGate supports DR by keeping secondary systems continuously synchronized with primary operational systems.

RTO and RPO — Recovery time objective (RTO) is the target time to restore service after an outage; recovery point objective (RPO) is the acceptable amount of data loss. GoldenGate is deployed in architectures where both must be kept low, including near-zero in financial systems.

Zero-Downtime Migration — A migration pattern in which old and new systems stay synchronized while users keep operating. GoldenGate enables online cutover by continuously replicating changes throughout the migration window, removing the maintenance outage.

ZeroETL — An architectural pattern that reduces or removes traditional batch ETL by continuously moving operational data into analytical or AI-ready environments. GoldenGate delivers ZeroETL-style patterns through real-time CDC and automated, schema-aware delivery.

Data Fabric — A managed architecture that connects data producers and consumers across systems, clouds, and domains. GoldenGate acts as the real-time data movement layer, and control plane, within an enterprise data fabric.

Real-Time Data Product — A curated, governed, continuously updated dataset or stream that can be consumed by applications, analytics platforms, AI systems, or business teams. GoldenGate produces real-time data products rather than simply moving raw rows.

GoldenGate Data Streams — A publish/subscribe interface for consuming GoldenGate change events. It lets applications, services, and AI agents subscribe to trusted, real-time operational data without a separate Kafka cluster.

AsyncAPI — An open specification for describing event-driven APIs. GoldenGate Data Streams exposes real-time change events to consumers through an AsyncAPI-based pattern.

GoldenGate Stream Analytics — Oracle's stream-processing capability associated with GoldenGate, supporting real-time processing, enrichment, analysis, and event-driven logic, including LLM and machine-learning operators, directly on streaming data.

GoldenGate AI Microservice — A GoldenGate component that brings AI-oriented processing into the pipeline, including LLM integration, embedding generation, enrichment, and classification, applied to change data in motion.

Vector Embedding — A numerical representation of data that lets AI systems compare semantic meaning. Embeddings underpin vector search, RAG, recommendation, and agent knowledge systems, and GoldenGate can generate them inline during replication.

RAG (Retrieval-Augmented Generation) — An AI pattern in which a model retrieves relevant external information before generating an answer. GoldenGate keeps RAG systems current by streaming operational changes into the knowledge bases and vector stores they draw on.

GraphRAG — An extension of RAG that uses graph relationships and connected context to improve retrieval and reasoning. Real-time operational data keeps the entities and relationships behind GraphRAG current.

Oracle AI Vector Search — An Oracle AI Database capability for storing and querying vector embeddings. GoldenGate supports real-time vector pipelines that keep vector search indexes aligned with operational changes as they happen.

Oracle AI Database — Oracle's AI-era database release, combining transactional data management with AI-native capabilities including vector search, JSON Relational Duality, in-database machine learning, and AI-integrated processing.

JSON Relational Duality — An Oracle AI Database capability that lets applications work with JSON documents while the database stores the data relationally. GoldenGate can capture changes from these structures and propagate them as complete, consistent business objects.

Veridata — Oracle's GoldenGate data-verification technology for comparing source and target data. It provides algorithmic proof of whether distributed systems are synchronized and supports auditability for compliance-sensitive environments.

Observability — The ability to monitor, understand, and troubleshoot system behavior. For GoldenGate, observability spans lag, throughput, process health, checkpoints, errors, and end-to-end pipeline performance.

Governed Data Movement — Moving data in accordance with security, compliance, access, lineage, and audit requirements. In AI systems this is critical, because agents and models must only reach data they are authorized to use.

Agentic AI — AI systems that take actions, use tools, trigger workflows, or interact with enterprise systems. Agentic AI depends on current data, governed permissions, auditability, and runtime control, the properties GoldenGate exists to provide.

Data Truth — The real-time, trusted data foundation that lets AI systems act on current enterprise reality rather than a delayed snapshot. In this paper, that foundation is the role GoldenGate plays through continuous, transactionally consistent streams from operational systems.

FAQ: Oracle GoldenGate

What is Oracle GoldenGate?

Oracle GoldenGate is Oracle's enterprise platform for real-time, transactionally consistent data movement. It captures committed changes from operational systems, preserves transaction order and integrity, and delivers those changes to databases, streams, lakehouses, analytics platforms, and AI systems across on-premises, cloud, multicloud, and edge environments.

Why does Oracle GoldenGate matter for enterprise AI?

Oracle GoldenGate matters for enterprise AI because AI depends on current, trusted business data. Models, RAG systems, and AI agents cannot reliably reason or act when working from stale batch snapshots or incomplete copies of operational data. GoldenGate streams business events as they happen, giving AI systems access to fresh operational reality while preserving governance, consistency, and auditability.

What are the main use cases for Oracle GoldenGate?

Oracle GoldenGate is used for mission-critical high availability and disaster recovery, zero-downtime migrations, real-time analytics, lakehouse ingestion, event-driven architectures, AI and RAG pipelines, data fabric architectures, and governed data synchronization across hybrid and multicloud environments.

How is Oracle GoldenGate different from traditional ETL?

Oracle GoldenGate differs from traditional ETL in how and when it moves data. Traditional ETL typically extracts data in scheduled batches and can place load on source systems. GoldenGate uses change data capture to stream committed database changes continuously from transaction logs, which removes dependency on batch windows, keeps downstream systems current, and avoids many of the operational risks of bulk extraction.

What is change data capture, and why is it important?

Change data capture, or CDC, is the process of identifying and moving only the data that has changed. CDC is important because it enables continuous synchronization without repeatedly copying entire datasets. For AI, analytics, and operational applications, CDC delivers fresher data with lower source-system impact than batch extraction.

What makes Oracle GoldenGate different from other CDC tools?

Oracle GoldenGate is built for mission-critical enterprise workloads where correctness, scale, availability, and operational trust matter. Its differentiators include deep Oracle Database integration, transactionally consistent delivery, broad heterogeneous platform support, mature failure recovery, support for high availability architectures, managed cloud deployment options, and capabilities for AI, lakehouse, streaming, and agentic workloads.

Is Oracle GoldenGate only for Oracle Database?

Oracle GoldenGate is not limited to Oracle Database. It has deep advantages with Oracle Database but is a heterogeneous data integration platform, supporting many Oracle and non-Oracle sources and targets, including relational databases, cloud services, Kafka-compatible systems, lakehouse formats, streaming platforms, and AI-oriented data architectures.

How does Oracle GoldenGate support RAG and AI agents?



Oracle GoldenGate supports RAG and AI agents by continuously delivering operational data into vector databases, RAG pipelines, AI knowledge bases, lakehouses, and event-driven systems. Using GoldenGate AI capabilities, changed records can be enriched, embedded, summarized, classified, or routed as part of the pipeline, so AI systems operate on current business context rather than stale snapshots.

What is the difference between Oracle GoldenGate and Kafka?

Oracle GoldenGate and Kafka solve related but different problems. GoldenGate specializes in high-fidelity capture and delivery of transactionally consistent change data from operational systems, while Kafka is commonly used as an event distribution platform. GoldenGate can publish to Kafka-compatible topics, and GoldenGate Data Streams can also provide a native publish/subscribe pattern for consumers that need trusted change events without a separate Kafka infrastructure.

What is OCI GoldenGate?

OCI GoldenGate is Oracle's fully managed cloud service for GoldenGate. It provides GoldenGate's capabilities without requiring customers to manage the underlying replication infrastructure, patching, scaling, and service operations themselves. It is intended for organizations that want enterprise-grade data movement with the operational simplicity of a managed cloud service.

When should an organization use Oracle GoldenGate?

Oracle GoldenGate is the right fit when data movement is mission-critical, when latency or transaction order matters, when source-system impact must be minimized, or when the architecture must support AI, analytics, high availability, migrations, or multicloud data products at enterprise scale. It is especially appropriate when downtime, data loss, duplicate transactions, stale AI context, or weak auditability would create business risk.

When might Oracle GoldenGate not be the right tool?

Oracle GoldenGate may be more capability than needed for simple low-volume batch ETL, small development sync jobs, or workloads where latency, transaction fidelity, availability, and governance are not important. In those cases, a simpler batch or low-cost integration tool may be sufficient. GoldenGate is designed for enterprise problems where data movement is part of critical business infrastructure.

How does Oracle GoldenGate help with governance and trust?

Oracle GoldenGate supports governed data movement by preserving transaction context, enabling observability, securing data in motion and at rest, and providing verification through GoldenGate Veridata. For AI and regulated workloads, this helps organizations answer whether data was delivered correctly, whether systems are synchronized, and whether downstream AI and analytics systems are acting on trustworthy information.

What is Oracle GoldenGate's role in a data fabric?

Oracle GoldenGate serves as the real-time movement and synchronization layer in a data fabric architecture, connecting data producers to data consumers. It captures operational events, moves them across systems, supports transformation and enrichment in motion, and delivers trusted data products to analytics, AI, application, and streaming environments.



Citations and References

- [1] A. Challapally, C. Pease, R. Raskar, and P. Chari, "The GenAI Divide: State of AI in Business 2025," MIT Project NANDA, July 2025. https://nanda.media.mit.edu/ai_report_2025.pdf
- [2] J. Ryseff, B. F. De Bruhl, and S. J. Newberry, "The Root Causes of Failure for Artificial Intelligence Projects and How They Can Succeed," RAND Corporation, RR-A2680-1, 2024. https://www.rand.org/pubs/research_reports/RR-A2680-1.html
- [3] MuleSoft (Salesforce), "2026 Connectivity Benchmark Report," 2026. <https://www.mulesoft.com/lp/reports/connectivity-benchmark>
- [4] Gartner, "Lack of AI-Ready Data Puts AI Projects at Risk," press release, February 26, 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-02-26-lack-of-ai-ready-data-puts-ai-projects-at-risk>
- [5] Oracle, "Oracle GoldenGate" (product overview), Oracle Corporation, 2026. <https://www.oracle.com/integration/goldengate/>
- [6] Oracle, "Oracle GoldenGate Data Sheet," Oracle Corporation, 2026. <https://www.oracle.com/a/ocom/docs/oracle-goldengate-datasheet.pdf>
- [7] National Institute of Standards and Technology, "FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard," August 13, 2024. <https://csrc.nist.gov/pubs/fips/203/final>
- [8] Oracle, "Regional Availability for Multicloud" (Oracle Database@AWS, @Azure, @Google Cloud), Oracle Cloud Infrastructure Documentation, 2026. <https://docs.oracle.com/en-us/iaas/Content/multicloud/regions.htm>
- [9] PwC, "PwC's 29th Annual Global CEO Survey," PwC, January 2026. <https://www.pwc.com/gx/en/issues/c-suite-insights/ceo-survey.html>
- [10] B. Jensen, D. Deano, M. Allison, and T. Bin Asad, "Bridging the Great AI Agent and ERP Divide to Unlock Value at Scale," McKinsey & Company, January 9, 2026. <https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/bridging-the-great-ai-agent-and-erp-divide-to-unlock-value-at-scale>

ORACLE CORPORATION

Worldwide Headquarters

Nashville, TN 37203 USA

Worldwide Inquiries

TELE + 1.650.506.7000 + 1.800.ORACLE1

FAX + 1.650.506.7200

oracle.com

CONNECT WITH US

Call +1.800.ORACLE1 or visit oracle.com. Outside North America, find your local office at oracle.com/contact.



blogs.oracle.com/oracle



facebook.com/oracle



x.com/oracle

Copyright © 2026, 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. 0626

Summer 2026

Author: Oracle Product Development



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

