Ovum
TMT intelligence | informa

The always-on business: Getting on the same page with Oracle GoldenGate

An Ovum white paper for Oracle

Publication Date: September 2016

Tony Baer, Principal Analyst, Information Management

Summary

Catalyst

Always-on, data-driven, and real time have become table stakes for enterprises in today's global economy. For instance, customers for retailing, digital entertainment, and gaming sites expect immediate engagement, where downtime is not an option. The same holds true for manufacturers optimizing supply chains, financial institutions guarding against intrusions, hospitals delivering patient care, and scheduled transportation carriers. All demand real-time visibility with the assurance that the information is current and in sync across all systems and locations. The variety of data sources, from transaction to messaging, document data, and the Internet of Things (IoT), is steadily growing, while deployment options are getting more distributed. With transaction volumes skyrocketing and data proliferating, the stakes for delivering a consistent picture in real time have never proved more urgent – and challenging. The challenges for getting a consistent, up-to-date picture are compounded as enterprises embrace cloud deployments.

Ovum view

Delivering a consistent, real-time view that is always in sync is not simply a technology challenge. It is a challenge that impacts how a business responds to inevitable disruptions in its market. It is a challenge for organizations to take cost out of their operations, execute faster, and harness new revenue opportunities. A flexible, balanced integration strategy is essential – there is a synergy between the bulk updates on which organizations have long relied and the use cases requiring low-latency, continuous updates that should be coordinated. The bulk loading of a data warehouse or Hadoop target may be triggered as the byproduct of a real-time event for synchronizing transaction systems. The change data capture logic should be interchangeable and adaptable between both scenarios to provide organizations with the necessary flexibility to adapt and respond. The solution must minimize overhead by capturing changes at the source and provide resiliency by isolating the metadata essential for preserving transactional integrity from the physical movement of data. And the solution must be platform-independent, capable of delivering consistent high performance beyond the SQL database "usual suspects" to the emerging array of NoSQL databases, Hadoop platforms, and streaming engines on which organizations are relying. As organizations embrace hybrid cloud deployment, the solution also must provide the assurance that applications operating from databases on- and off-premise deliver consistent, reliable results.

Key messages

- Real time is driving service-level expectations, even for offline workloads.
- With enterprises embracing hybrid cloud deployment, the need to keep data in sync, in real time between cloud and on-premise systems becomes table stakes.
- Real-time data integration is not simply a technology argument; it is essential to the ability to execute faster, which can lead to bottom-line cost reductions and top-line creation of new business and revenue opportunities.
- Supporting real-time business requires a data-replication technology that is flexible, scalable, performant, and platform-independent, with little to no overhead.

Enterprises must prepare for real time

Figure 1: Real-time is driving SLA requirements

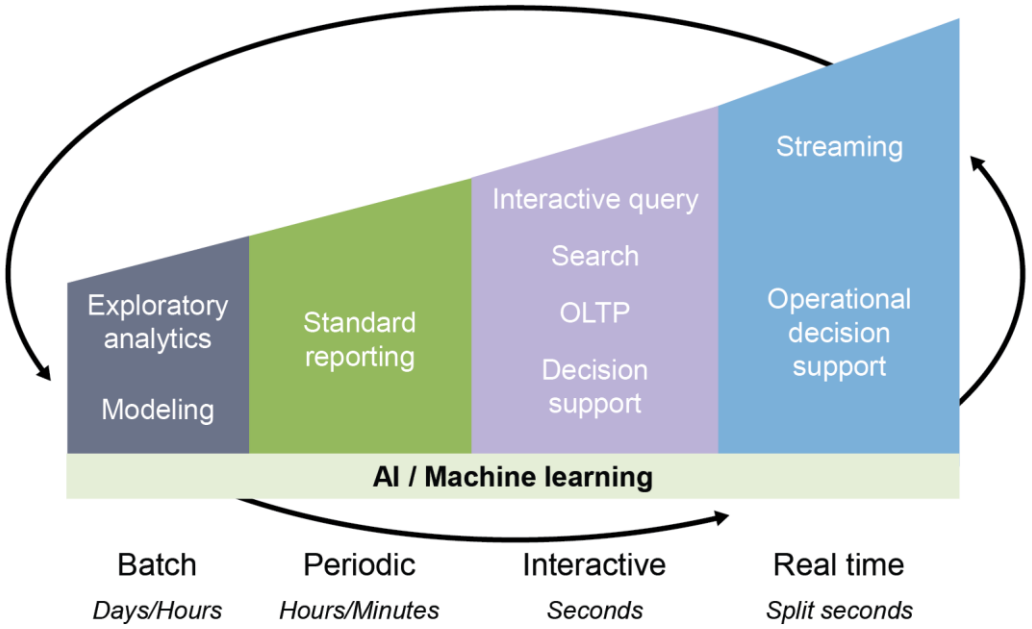

Source: Ovum

Always live, always on the same page are table stakes

In an increasingly interconnected world, global enterprises are expected to provide a consistent picture. They cannot provide different answers regarding product or service availability or delivery promises to consumers. Business-to-business organizations cannot afford to give conflicting inventory data to supply chain partners. Financial services companies cannot afford to show discrepancies in trading data. And as enterprises embrace the cloud, they cannot tolerate discrepancies between on-premise and cloud-based applications. For these organizations, real-time replication and disaster recovery provisions are no longer luxuries – they are essential to keeping the heartbeat of their businesses alive while improving their ability to execute.

Being on the same page is about keeping the pulse of the business alive. That involves integrating real-time replication as a continuing process, ensuring delivery of service and seamless disaster recovery. The recent outage of the global reservation system of a major US airline during the height of the summer travel season – after a cascade of events beginning with a power interruption and the failure of routing equipment – revealed the drawbacks of failing to build real-time backup and replication into routine operations. The failure grounded thousands of flights, a disruption that took several days for the airline to recover from, causing significant revenue and reputational loss.

The applications and use cases required for running a global enterprise have always varied; increasingly, service expectations, even for offline batch workloads, are being driven by real-time

business, regulatory, and operational challenges, as shown in Figure 1. What has changed is the gravity toward engagement in real time, whether that involves

- interacting with the customer
- digesting the growing volume of device and IoT data enabling the organization to transform responsibilities such as risk mitigation;
- intrusion protection, and/or
- manufacturing or supply chain optimization from cost centers to competitive differentiators.

"Right time" – Batch and real time are no longer worlds apart

Batch, periodic, interactive, and real-time workloads have historically operated separately; however, the increasingly connected environments in which global enterprises operate are driving the need for "right time" integration. A real-time event, such as the sudden reaction of fans when a sports team scores an upset, a disruption to the supply chain, and/or the need to respond to a real-time intrusion, will increasingly drive operations that are also conducted offline. For data integration and updating, that drives demand for a multipronged, right-time strategy that choreographs fast bulk-data transfers with continuous real time; an enterprise will require the right blend of both. Consider the following scenarios.

Use Cases

The cloud beckons…

It is getting impossible to ignore the growing enterprise embrace of cloud computing. Ovum forecasts that enterprise cloud services spending will grow at a compound annual growth rate of 22% through 2020. As enterprises embrace the cloud, data gets more dispersed, and the urgency for keeping on-premise and cloud-based applications in sync grows more urgent.

There are many drivers behind the move to cloud. Data is getting more dispersed. Smart mobile devices and the Internet of Things are generating data that originates outside the four walls of the enterprise and naturally lives in the cloud. Enterprises are also looking to the cloud as the place for rapidly testing and spinning up new applications and services without disrupting legacy production systems, while others are seeking to shift expenses from capital to operations.

For many organizations, the cloud enables them to focus on their core competency of delivering their own products and services, rather than focusing on building and maintaining IT infrastructure. That was the driver with a major software-as-a-service provider of accounting programs for small businesses (SMEs). It currently hosts its product, which is based on nearly 30 clusters of redundant Oracle 11g and 12c databases that run OLTP and analytic workloads, delivering five nines uptime to its 1 million-plus SME client base. In planning for future growth, the provider is facing a decision that is very similar to that of most enterprises with their own data centers: Is its core competency running a data center or delivering value through software? For the SaaS provider, the answer is a bit ironic: Its core IP is delivering business accounting software, not cloud services per se. So it has decided to get out of the hosting business and deliver its software through a third-party public cloud.

SaaS provider relies on bidirectional, active-active replication

To avoid disrupting its client base, it is migrating its accounting, payments, and payroll applications on a phased basis over 18–24 months. That will require side-by-side operation of both sites during the migration period. For the phased migration, it requires a solution that supports side-by-side operation by keeping existing systems in its data center in sync, in real time, with those of its new cloud provider. The provider was already an Oracle customer, using Active Data Guard for backup and disaster recovery. With the migration, that needed to shift to a real-time, bidirectional, active-active replication solution. It chose Oracle GoldenGate for two reasons. First is performance, which allows real-time change data capture between data centers. Additionally, it chose GoldenGate for its multiplatform support, which supports replicating with non-Oracle sources and targets, such as its HP Vertica data warehouse and Hadoop big data analytics platforms, which are used for delivering enterprise performance management services to clients.

Brick-and-mortar retailing – real-time inventory key to survival

Under the onslaught of online retailers such as Amazon and eBay, traditional retailers are fighting for survival – aiming to combine the best of both worlds in leveraging the personal experience in the store and the immediacy and broader availability of SKUs online. For an 800-store national retailer, optimizing inventory is part of the secret sauce for combining the best of the brick-and-mortar and online worlds. For instance, if Texas has a freak cold snap but most of the inventory for winter coats sits in Maine, the retailer must determine how to juggle inventory from closer locations to meet demand where it is occurring. Turning inventory on a dime requires the reorchestration of central inventory, regional inventory, logistics, and point-of-sale systems – and all must read from the same sheet of music.

This has cascading effects, not only at the point of sale in the store, but also with programs such as regional, location, and/or customer segmentation-based promotions; the retailer obviously cannot afford to allow promotions to fall out of sync with inventory. Likewise, if the customer is in the store and needs to order an item that is on back order or not stocked locally, the store must be able to provide firm delivery-date promises while the customer is at the counter. The technology challenge for this retailer is that inventory and point of sale are based on different systems; central inventory, which drives supply chain movements to distribution centers, online fulfillment, and retail stores, is a legacy application based on IBM DB2, while the point-of-sale systems, which are used for decrementing in-store inventories, are on Oracle-based systems.

Real-time change data capture keeps point of sale in step with central inventory

The retailer uses Oracle GoldenGate to keep the Oracle-based point-of-sale systems in sync with DB2 central inventory with subsecond latency – quick enough to provide customers with the right answer when they are at the website or point of sale, or to provide inventory planners with the real-time snapshot they need in case sudden events, like the upset win of a sports team, require rapid juggling of merchandise across the supply chain.

Rationalizing data pipelines across major banking institution

Banking institutions have historically had to manage multiple transaction, analytics, and archiving silos as a legacy of having dedicated systems for specific products, such as checking, savings, loans, and credit cards, along with residential and commercial lending. Yet banking customers take for granted

that, when they log in to online banking, the information is available and up to date on all of the accounts they maintain. For retail banks, this is no small feat.

To ensure a consistent picture across all target systems, the bank required a real-time, active-active data replication solution with native support for multiple platforms and emerging technologies. And with the bank actively upgrading, consolidating, and moving systems between data centers, and to the cloud, it required a common fabric that could ensure reliable, real-time delivery that supported multiple platforms. It needed a solution that could integrate with emerging open source technologies, such as Apache Kafka, that provides a highly scalable and economical messaging platform that is useful for scenarios where you have a "firehose" of events, which has become a de facto standard in the big data community.

Performance and platform independence key to ensuring reliable customer service

After testing well over a dozen alternatives, the bank chose Oracle GoldenGate because of its performance and platform independence. By relying on GoldenGate, the bank could standardize on a common tool that provides the fabric for moving data from any source to any target, with low overhead and high performance that ensures that active systems across the bank provide a common, consistent picture.

Among the bank's requirements, it has strict data-retention policies for transaction systems, keeping data for 125 days before pushing it off to archive. It uses GoldenGate to perform data movements on a nightly basis. To reduce reporting overhead on production systems, it uses GoldenGate's Big Data Adapter to perform change data capture live, feeding it to Apache Kafka messaging queues that populate Cassandra NoSQL databases that are used for Tableau self-service reporting. And, like many enterprises, the bank is embracing a hybrid cloud strategy, using GoldenGate for staged database migration to and side-by-side operation with the Amazon Web Services cloud.

What is the value of real-time, end-to-end data integration?

Avoided costs and business resiliency

Effective real-time, end-to-end integration through ongoing change data capture (CDC) can take cost out of operations by

- avoiding the time- and resource-intensive processes for reconciling data from systems spread across multiple geographies
- eliminating the need to reverse decisions and repeat transactions based on data that is out of date
- reducing the time and cost of disaster recovery, and virtually eliminating the downtime for system restoration.

These benefits have clear cost advantages, from savings of time and labor. Avoided downtime also provides intangible cost benefits, such as the revenue that can be gained because the business is always open. The cost of downtime has been an open issue ever since the dawn of computing. But as

recent events in the airline industry have proved, the cost of a single day of downtime for a major national carrier can result in the cancellation of thousands of flights.

Faster business execution

Operating with the assurance of up-to-date data means that the business can make smarter, actionable decisions faster, without the need to wait for the data to get refreshed. It allows faster execution. The retailer makes better decisions on where to stage inventory to fulfill customer orders, faster. The bank gains a better understanding of the financial positions of its customer base and can more effectively target new offers such as credit lines. A global online music and entertainment provider can deliver more reliable service when content is properly staged.

Ability to generate new revenue streams

IoT is presenting the opportunity for enterprises to add new value-added products and services and, in some cases, transform the business. For instance, with generational changes in car ownership, IoT is enabling automakers to reposition themselves as "transportation service providers." With smarter machinery, industrial equipment suppliers are creating new remote monitoring and optimization services for helping manufacturers increase uptime while reducing maintenance and operational costs; given the thin margins in a highly competitive global manufacturing sector, the ability to improve both the top and bottom lines matter when it comes to business survival. And with hot gaming fads such as Pokémon Go driving more traffic to smartphone and consumer electronics retailers, Bluetooth sensors are enabling retailers to monitor traffic and direct instant, location-based promotions to customers.

The common thread for all of these IoT business examples is that they must execute in real time and require current data that changes by the second – or less. So databases that track customers' products must be kept in sync with the latest real-time events. Updating of all data sources must be an ongoing process.

The new technology environment: Variety *is* the spice of life

"Polyglot" IT environments becoming the norm

Ever since midrange and distributed systems ended the monopoly of the mainframe, the ideals of harmonizing platforms have waxed and waned. With the emergence of open systems, relational databases became the de facto enterprise standard. But growing bandwidth, smart sensors, social media, and the mobile revolution unleashed vast torrents of new forms of data. Beyond transactions, the addition of clickstream and other log data, IoT sensor and device data, geolocation data, and complex JSON document data had to be accommodated. This drove demand for a new generation of data platforms that could handle distributed operational processes and big data analytics. This led to an outpouring of new data platforms that were purpose-built for these new operational and analytic workloads, as shown in Figure 2. Most large enterprises today have a mix of these platforms.

While the new breed of data platforms is now adding overlapping capabilities – such as SQL relational databases handling JSON data, or Hadoop supporting interactive SQL query – today's global enterprises are more likely than not to rely on an ecosystem of multiple platforms, each optimized for specific mixes of workloads. In these environments, there is no single source of the truth. For instance, customer data is likely to be spread across these data silos, with transaction data stored on enterprise relational databases, customer profile data recorded on NoSQL platforms, and whole log files pertaining to the customer's navigation of your organization's online site residing on Hadoop.

Figure 2: The new data processing silos

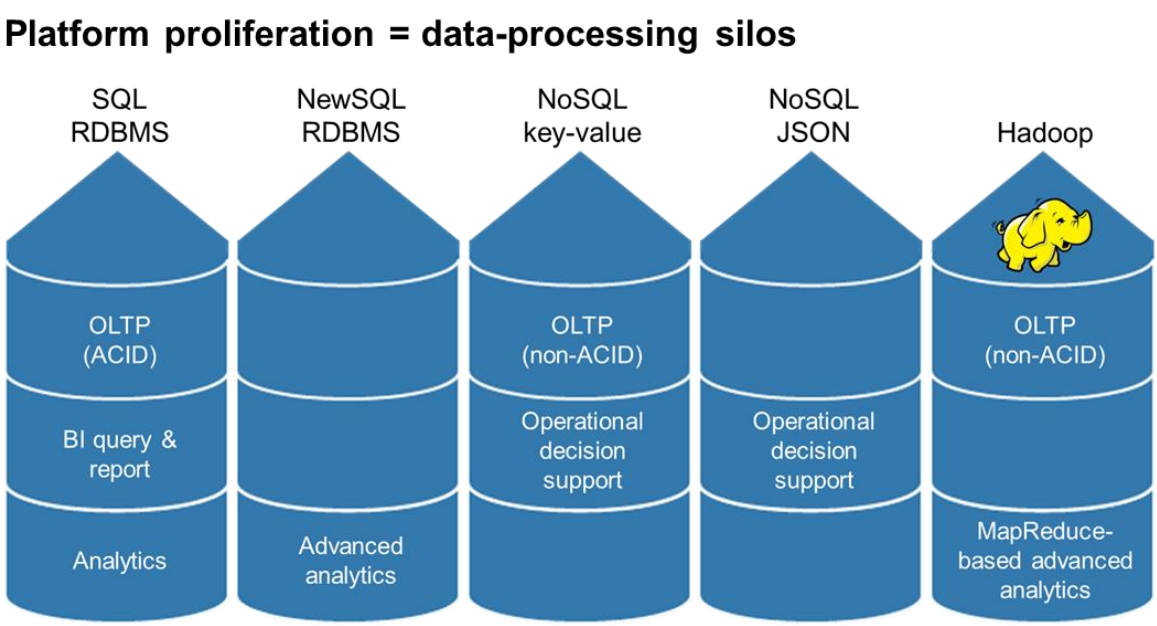

Source: Ovum

Hybrid cloud is the new norm

Part of the polyglot environment is the emergence of cloud alongside on-premise systems. Cloud adoption is spreading from test/development and new workloads to core enterprise applications. Ovum's enterprise insights data reveals that cloud investments are going mainstream, with nearly half of all enterprises globally reporting new or major transformation investments in the cloud during 2016 (see Figure 3).

Real time is driving SLA expectations across all workloads

As noted earlier, behind the scenes at every "real-time" business is a wide spectrum of applications and use cases that help the organization understand, transact, report, and respond to phenomena that are both long-running and, increasingly, immediate.

That remains the case; for instance, modeling and segmentation processes are intended to be thorough and therefore conducted offline. But increasingly, these offline processes are being driven by real-time events, requiring the ability to adapt and learn from changes in the competitive and operational landscape. Unexpected surges in customer demand, as detected through the social media firehose, may prompt a reappraisal of existing customer segmentation models. A unexpected

disruption in the supply chain, the recall of an automotive part, or the sudden detection of an intrusion or data breach are all events that impact not only an organization's immediate response to events, but also fundamental assumptions of an organization's underlying business, operational, and/or regulatory compliance strategy. It may be necessary to shift the change data capture of a specific cohort of data from fast bulk to real time, either temporarily to deal with a sudden disruption or permanently to accommodate a structural change in the competitive or operational landscape.

Furthermore, data replication strategies should also enable the organization to take advantage of new technologies, such as the Spark computing engine, which has introduced real-time processing to big data analytics, making practical the application of advanced machine learning and streaming that can provide organizations with a competitive edge on how they adapt and respond to changing markets. As a fast-emerging big data compute engine, Spark is already transforming traditional offline processes such as modeling and segmentation to become more adaptive in real time.

To enable the organization to adapt to change and leverage new, game-changing analytic approaches, the change data capture logic must be portable between bulk and real-time tooling or processes so the shift from offline to real time is seamless.

Figure 3: 2016 cloud spending trends 2016

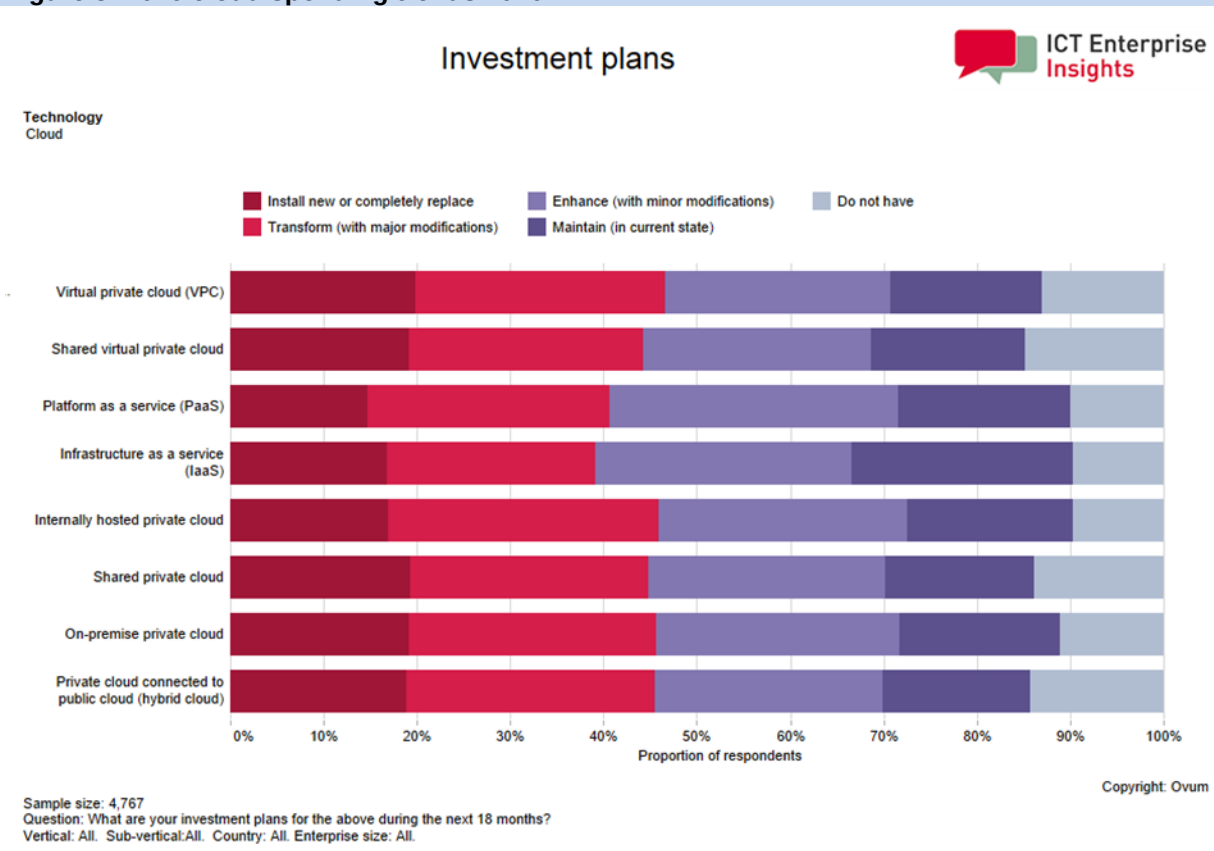

Source: Ovum

While expectations for SLAs are changing, the common thread remains that the data must be current. But for offline systems that perform functions such as modeling, fast bulk solutions are the most practical answer. Given the wide range of use cases and service levels, data integration and data movement solutions should play well together; the change data models should be interoperable so

data movement processes adhere to common logic, and the data targets are updated or populated consistently.

What's essential for real-time, end-to-end data integration?

Performance and scalability

Among the core ingredients for delivering end-to-end real-time data integration is high performance. Solutions must be capable of low-latency, *continuous* capture and delivery of changed data from source to target with minimal overhead and the ability to support high-volume throughput. To combine performance and scalability, the architecture should minimize overhead – ideally by abstracting transaction protection from physical capture and delivery – and support highly parallel execution.

Reliability

The data must be current and valid. Solutions must ensure ACID transactional integrity from source to target. This is especially critical for organizations that operate across multiple data centers or are operating in side-by-side mode as they migrate on-premise systems to cloud. Transactions must be preserved and kept in sequence, maintaining the reliability, accuracy, and sequence of transaction data, with full referential integrity.

Heterogeneity

Because polyglot IT environments are the norm, the solution must be fully platform-independent. It must be capable of capturing and delivering data between all data platforms that are active throughout the enterprise. This encompasses all major relational database platforms, NoSQL databases, Hadoop, and cloud platform targets, along with message queuing pipelines feeding streaming data.

Always on

Downtime is not an option for today's global enterprises – businesses cannot afford to take their customer-facing operational systems offline for traditional batch updates, planned maintenance, or unplanned downtime. As noted above, low overhead is essential for the system to operate while production is live. And for organizations with footprints spanning multiple geographies, high-availability disaster recovery capabilities are essential.

How Oracle GoldenGate supports the always-on business

ACID support, performance, and resiliency

Oracle GoldenGate is designed for continuous ACID-compliant, low-latency, real-time CDC that can serve one or more targets thanks to a unique, modular architecture that decouples capture and

delivery steps from canonical trail files (see Figure 4). This versatility, performance, and resiliency is not matched by most niche vendors.

GoldenGate captures metadata critical to safeguarding ACID compliance through canonical trail files that are kept separate from the actual capture and delivery operations; this ensures that transaction properties and sequence are maintained regardless of what occurs during the replication process, ensuring reliability and recoverability. Its change log capture was designed to impose minimal overhead on the source database, and by keeping metadata separate, makes capture and delivery a sufficiently lightweight process that it can provide subsecond performance from source to target – a level of performance and reliability matched by few if any niche tools or open source technologies. Likewise, this architecture allows much faster recoverability compared to most niche tools.

Figure 4: Oracle GoldenGate architecture

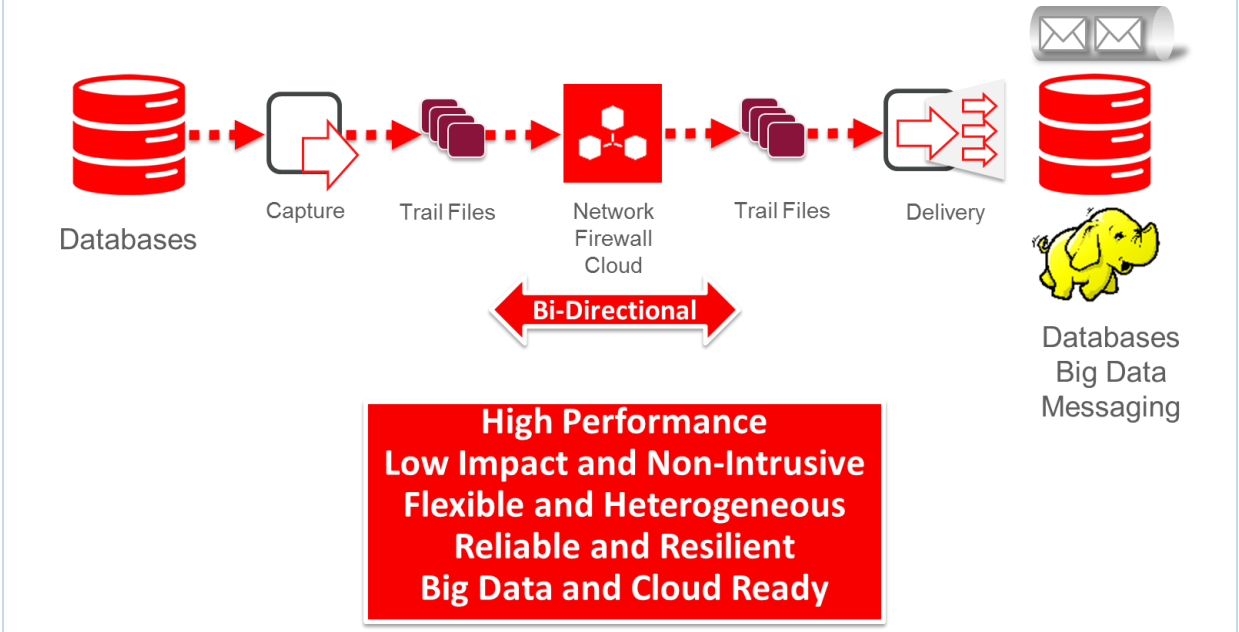

Source: Oracle

Flexibility

Supports multiple modes

Because of its modular architecture, Oracle GoldenGate operates with a flexible topology that allows a greater variety of replication modes than many niche CDC and data replication tools. That allows GoldenGate to support both homogeneous and heterogeneous architectures, making it adaptable to the multidata platform environments that are becoming more common in today's organizations. And it can operate closely with Oracle Data Integrator by sharing common CDC models, providing an end-to-end solution that accommodates the variety of real-time and bulk-data-loading scenarios that are common to global enterprises with diverse IT environments and applications with varying SLA requirements.

Options range from unidirectional or bidirectional, point-to-point replication to active-active multimaster mode, peer-to-peer, broadcast, and consolidation modes. And by eliminating the need for a central

replication hub, GoldenGate reduces overhead and eliminates single-point-of-failure vulnerabilities. As a result, GoldenGate supports a wide variety of use cases, encompassing

- query offloading, operational reporting, and zero-downtime data migration
- publish/subscribe scenarios
- data center or private cloud consolidation
- data distribution via capture and/or delivery to messaging sources/targets
- fully active distributed database operation for high availability and zero-downtime disaster recovery
- event-driven applications
- streaming data support for real-time data warehousing and massive parallelization.

Designed for the cloud

Oracle GoldenGate and Oracle Data Integrator (ODI) complement each other in supporting migrations to public or private clouds. ODI provides the fast bulk-data movement to populate the cloud, while Oracle GoldenGate complements it by providing real-time change data capture updates. GoldenGate leverages the same change data movement models and configurations developed in ODI, avoiding the need to reinvent the wheel. Clearly, security is critical when migrating to the public cloud; both ODI and Oracle GoldenGate support a variety of security options, including secured networks, restricted networks (where only specific ports are authorized and the cloud IP target is whitelisted), and restricted network access through HTTP proxy (where whitelists do not exist).

Database independent

While Oracle GoldenGate 12c integrates directly with the multithreaded ingest capabilities of Oracle Database 11g and 12c, the tool was designed to be database independent, supporting non-Oracle databases and messaging targets through ODBC and JMS. For non-Oracle databases, GoldenGate offers a coordinated delivery feature that avoids the need to manually manage and configure multiple parameter files.

Extending to big data

The big data world has been the domain of emerging open source technologies that maximized the power and scale of low-cost, commodity infrastructure. Traditionally limited to batch processing, in recent years the Hadoop platform has matured significantly with its ability to run multiple types of workloads, such as interactive SQL and Spark. Governance has improved, with the ability to selectively encrypt data in motion and at rest. But the capabilities of open source technologies for providing change data capture – real time or not – are still lacking.

Oracle GoldenGate for big data is among the first to deliver subsecond performance and resiliency for replication to the emerging platforms and technologies in the big data world, including Hadoop, NoSQL databases, and streaming systems (see Figure 5). It complements open source technologies – integrating with them and picking up where they leave off. Oracle GoldenGate's subsecond CDC capability, coupled with the real-time computing of Spark, provides a compelling combo that could integrate the predictive and prescriptive analytics of Hadoop with the real-time operation of enterprise OLTP systems. As transaction systems operate, Hadoop can instantly pick up the latest updates and

take advantage of the Spark compute engine's portfolio of machine learning, graph, and streaming analytics capabilities. For instance, with updates to the supply chain management system, Hadoop can apply machine learning side by side to uncover hidden trends that could predict where demand is headed. Supply chain-optimization applications could generate prescriptive or next-best-action recommendations based on the generated predictions.

GoldenGate's support of the big data technology ecosystem includes bidirectional integration with Apache Kafka for feeding and ingesting from real-time data feeds. Additionally, there is support for Flume, Hadoop's original high-speed data ingest engine, and the Apache Storm streaming engine. By supporting Kafka and Flume, GoldenGate provides a much lighter-weight, less-invasive approach to populating Hadoop from SQL databases compared to Hadoop's original Sqoop utility. And thanks to GoldenGate's declarative design, creating data mappings to these pipelines is greatly simplified. GoldenGate populates Hadoop's storage and data management engines, including HDFS, Hive, and HBase; there is also support for feeding JSON data to NoSQL databases such as MongoDB.

Figure 5: Oracle GoldenGate for big data

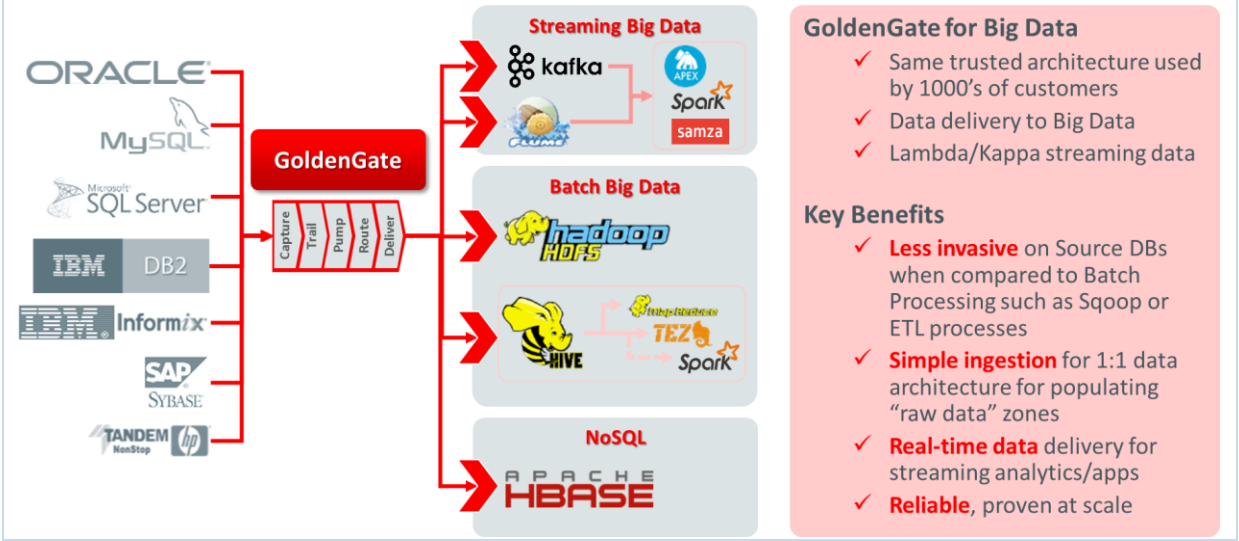

Source: Oracle

Takeaways

Being always on has become table stakes for global enterprises, whose customers expect real-time engagement. The key to becoming an always-on enterprise is ensuring that all the diverse data platforms that are maintained across the organization – often in different geographies – are kept in sync. In so doing, organizations cannot afford to lose resiliency; transactions cannot be lost, corrupted, or recorded out of sequence in the process.

This requires an end-to-end strategy that encompasses fast bulk-data transport for updating targets such as enterprise data warehouses with the continuous, real-time updates that are required of heartbeat systems that run the enterprise and ingest streaming data feeds from sources such as the Internet of Things. Change data capture logic must be interchangeable and adaptable between both scenarios to provide organizations with the necessary flexibility to shift different process between bulk and real time as competitive requirements dictate.

Unlike niche CDC tool providers, Oracle GoldenGate is differentiated by its unique approach to resiliency by separating the trail (metadata) from the physical capture and transmission process – providing critical protection of transaction integrity. It also differentiates with support for interoperation with Oracle Data Integrator, allowing sharing of CDC logic and scripts between fast bulk transport and real time. And Oracle GoldenGate differs from open source loaders to Hadoop, such as Apache Sqoop and Apache Flume, by combining change data capture with real-time processing. With the connector to Apache Kafka, Oracle GoldenGate offers the best of both worlds – providing a feed with full transactional integrity to the Apache open source project that is fast becoming the de facto standard for message queuing to big data platform targets.

Appendix

Methodology

This white paper was produced based on product research, briefings with Oracle, and detailed discussions with Oracle GoldenGate clients.

Author

Tony Baer, Principal Analyst, Information Management

tony.baer@ovum.com

Ovum Consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Ovum's consulting team may be able to help you. For more information about Ovum's consulting capabilities, please contact us directly at consulting@ovum.com.

Copyright notice and disclaimer

Any views and/or opinions expressed in this product by individual authors or contributors are their personal views and/or opinions and do not necessarily reflect the views and/or opinions of Informa Telecoms and Media Limited.

CONTACT US

www.ovum.com
analystsupport@ovum.com

INTERNATIONAL OFFICES

Beijing
Dubai
Hong Kong
Hyderabad
Johannesburg
London
Melbourne
New York
San Francisco
Sao Paulo
Tokyo