Oracle extends Autonomous Database to transaction processing

Publication Date: 03 Aug 2018 | Product code: INT002-000155

Tony Baer

Ovum view

Summary

Oracle has taken the next step in the lights-out cloud database with the general release of Oracle Autonomous Transaction Processing (ATP) Database 18c. In contrast to the initial release, which addressed data warehousing, Oracle ATP will offer more control of autonomy. This reflects the fact that in OLTP, applications drive the database and typically carry their own optimizations that might otherwise conflict with Oracle's default settings. ATP addresses a more complex use case than data warehousing, as the workloads are more mixed (with some analytics often thrown in). The workloads are also typically more mission critical because they directly support business operations. Ovum believes, therefore, that the learning curve for optimizing ATP will be longer compared to that for data warehousing, because of the sheer variability. But, in the long run, we expect that autonomy (defined as machine learning adding optimization atop database automation) will become the norm with *managed* cloud database services (i.e. database-as-a-service or DBaaS). Oracle is the first database provider to go down this path. Therefore, the release of ATP provides an opportunity for enterprises with DBaaS plans to start kicking the tires and testing to find the right level of autonomous operation that will work with the way their applications already run. The good news is that Oracle customers won't have to change their indexes to run on ATP. More to the point, ATP provides an early opportunity for enterprises considering DBaaS deployments to evaluate how databases run with machine learning (ML) can help them save money.

The next step in Oracle's autonomous database journey

When Oracle took the first step last fall with release of the Autonomous Data Warehouse (ADW), it indicated that OLTP and NoSQL databases would be next on the list (along with middleware) in 2018. The next step – ATP – has now gone into general release.

Like ADW, Oracle ATP is available only on the Oracle Public Cloud, where Oracle can control the environment. Oracle ATP is designed to fulfill the needs of demanding mission-critical databases and includes a scale-out Real Application Clusters (RAC) architecture running on fault-tolerant Exadata cloud infrastructure which provides application-transparent patching, upgrades, and failover. It is available with the same two service levels offered for ADW: an enterprise service that delivers three nines of availability (99.95%) and a mission-critical service that offers four nines (99.995%), which amounts to a half hour of downtime per year. These are among the highest SLAs offered for managed cloud operational database services, especially when exclusions (such as for maintenance, or service interruptions under five minutes, by Amazon and Google Cloud, respectively) are considered.

The mission-critical service adds Oracle Active Data Guard for high availability and disaster recovery to meet the stricter SLA. As an autonomous database, ATP will automatically run and optimize itself based on service levels defined by the user; automatically secure the data; and automatically repair and upgrade. The big difference with database automation is that ML is applied to optimize all the settings; diagnose errors; and determine how to repair, patch, or upgrade the system while maintaining SLAs and avoiding outages.

Oracle Autonomous Database uses a serverless architecture to support rapid elasticity for both compute and storage, allowing users to scale up as demand increases and scale down to reduce costs when demand decreases.

What's in it for the enterprise?

As we noted in the report *Oracle's Autonomous Database 18c: An Initial Assessment*, autonomous operation has been a journey built on years of developing automation for the different tasks of running a database. For instance, Oracle's databases have automated features such as storage management, memory management, statistics gathering, and query rewrites since the 9i generation over a decade ago. Among enterprise databases, Oracle is among the most automated. But other databases also have degrees of automation. For instance, Microsoft SQL Server provides agents that can automate certain security, event-handling, and error-tracking tasks. DBaaS cloud services typically take automation to the point where all deployment, patching, and configuration are handled without human intervention.

Oracle takes automation further by using Application Continuity and rolling upgrades on RAC to perform all maintenance online, delivering the availability that 24×7 OLTP applications need. Enterprises that are embracing DBaaS already accept that there is no value-add in having their DBAs worry about patching. Adding ML won't eliminate the DBA – they will still be necessary for more strategic tasks such as designing the database architecture; specifying the schema; defining security and SLA requirements; and conducting application-related tuning. And, as we note below, although Oracle is taking away a lot of the "knobs," it won't be able to take away all of them.

It is clearly early days for autonomous databases – Oracle is the first to offer them. But Ovum believes that applying ML to databases will eventually become the norm, allowing enterprises to more effectively use their DBAs and get the most out of their databases. However, OLTP is a more complex challenge than data warehousing because there are more moving parts. Although part of that is attributable to the nature of the workloads, the big factor is applications: the last thing that any enterprise (or Oracle) wants is to have their order-processing or accounting systems crash because the database implementation is incompatible.

Our advice to enterprises is that now is the time to begin kicking the tires and seeing how your applications will work with Oracle's autonomous features. It may be early, but as Ovum believes that ML will become the norm in operating databases, it makes sense for enterprises considering DBaaS deployments to get their feet wet in – literally – testing the waters.

Optimizations for OLTP

There are distinct differences in optimizing a database for OLTP as opposed to analytics. The use cases are obviously quite different: although many organizations are operationalizing analytics, the stakes for maintaining availability in OLTP are much higher because transaction processing is the heartbeat of the enterprise. Errors such as incorrectly applying a security patch or forgetting to apply one can bring the day-to-day operation of an enterprise to a standstill. Conversely, if analytics go offline, the organization should theoretically be able to continue operating.

There are physical differences between analytics and OLTP. OLTP databases are typically row based because of the need to access individual records, while increasingly analytics are relying on columnar tables. Furthermore, indexes that are key to fast lookups are more critical for OLTP than for analytics (which in some cases eliminate them). ATP uses ML to continuously analyze the workload and create new indexes online when needed. Beyond indexes, fast access to data in OLTP is through placing result sets in memory, whereas for analytics, it will be for caching data summaries.

Oracle ATP has built additional optimizations for query planning by changing the way statistics are collected for factors such as table size, distribution of values, and frequency of access. Normally, these stats are updated only periodically; in ATP, they are updated continually in real time and are used for testing old versus revised query plans.

Why Oracle ATP offers indexing "knobs"

Compared to data warehousing, online OLTP has more moving parts. For starters, the workloads are often more heterogenous, as many OLTP deployments also incorporate some degree of batch and analytics (typically query and reporting and in some cases ad hoc analytics). Furthermore, the current generation of enterprise applications often embed analytics into their transaction processes. For example, when you issue a procurement, there are typically capabilities to perform some what-if queries regarding supplier choice, quantity, timing, or other parameters. Furthermore, OLTP deployments often (although not always) require higher concurrency levels, as the user population is generally larger.

Then there is one other variable that complicates matters: applications. Compared to analytics, transaction applications are more likely to tap database features such as stored procedures that place some of the logic in the database tier. But Oracle determined that the biggest gain with optimizing (and automating) transaction databases was attributable to indexes. In a perfect world, the ML features of the database would decide which indexes to run, period. The catch, however, is the world is not perfect: enterprises have often spent years creating and tuning their indexing, and the last thing they want is disruptions from changing them.

So, Oracle is providing a couple options here. As a baseline, both options support running the indexes that enterprises have already created for their applications. The two options are

- run the database *only* with those original application-created indexes
- allow ATP to add or drop *additional* indexes that its ML features deem optimal for running the database.

As we noted earlier, ATP will be a learning experience for Oracle and its customers, and we expect that future releases will further optimize how applications of different configurations work with autonomous databases. It's a sensible starting point for Oracle on its ATP journey: database autonomy won't start by requiring Oracle customers to rewrite their applications.

Appendix

Further reading

Oracle's Autonomous Database 18c: An Initial Assessment, INT002-000011 (November 2017)

Oracle SaaS and PaaS Cloud: Assessing the Progress, IT0014-003178 (December 2016)

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

CONTACT US

ovum.informa.com

askananalyst@ovum.com

INTERNATIONAL OFFICES

Beijing

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

San Francisco

Sao Paulo

Tokyo