

IDC PERSPECTIVE

Oracle's Autonomous Database: AI-Based Automation for Database Management and Operations

Carl W. Olofson

David Schubmehl

EXECUTIVE SNAPSHOT

FIGURE 1

Executive Snapshot: Oracle Autonomous Database

This IDC Perspective examines the announcement and initial launch of Oracle Autonomous Database. It also looks at how a search vendor is planning to use Oracle Autonomous Database in its business operations and why.

Key Takeaways

- Oracle has recently announced its latest database product, a self-managing autonomous database.
- Oracle is at the vanguard of a movement to make enterprise software easier to deploy, use, and administer, using artificial intelligence and machine learning to provide automation capability requiring little or no human intervention to manage the software.
- While Oracle Autonomous Database is one of the first products in this new automated category, it won't be the last. Other vendors are building similar tools for third-party management of existing enterprise software packages, and many software vendors will also add this kind of capability to their offerings.

Recommended Actions

- Actively evaluate tools such as Oracle Autonomous Database for suitability in domains such as process automation, digital transformation, and other business processes that require database access. If this seems to be a fit, organizations should begin evaluating Oracle Autonomous Database now, rather than waiting and seeing how other organizations are faring.
- Start with low-risk pilots to identify the best opportunities for full-scale implementation.
- Develop metrics/measurement practices to determine how and where these systems are succeeding.

Source: IDC, 2018

SITUATION OVERVIEW

Introduction

Oracle recently announced Oracle Autonomous Database, an Oracle Cloud DBaaS offering based on Oracle Database, but with special features that enable it to self-tune and self-heal. This service relieves database administrators (DBAs) of the remaining operational tasks that normally accompany even a cloud-based database deployment. Those tasks include advanced tuning functions, database security, and troubleshooting. DBAs will soon be able to focus more time on design and development activities instead of spending time administering their database installation. Database administration will be fully automated, providing Oracle with the opportunity to guarantee 99.995% availability (including unplanned maintenance, bug fixes, configuration changes, and security patches) and no more than 30 minutes of downtime a year. The product will be tuned continuously, using machine learning (ML) models of observed behavior as patterns of user and system's behavior change.

Collectively, through its primary research, IDC has found that as much as 75% of the total cost of database management (including software and hardware costs amortized over five years) can, in some cases, be labor. Any DBaaS eliminates a large amount of the operational work (running backups, applying patches, etc.), but this capability redirects most of the remaining labor from "drudge" work to high-value work that directly benefits the enterprise. In the first version of this capability, Autonomous Database (version 18c) will only address analytic (that is, non-transactional) databases. However, the plan is to support all Oracle Database workloads over the course of the next year.

The Challenges of Managing a Database: Tuning, Troubleshooting, and Patching

Databases used for mission-critical activities, or for enterprise data analysis, present special challenges. Mission-critical transactional databases must not only be always available, but they must deliver excellent transaction throughput. Databases that support enterprise data warehouses can become massive in size, supporting very complex queries with many nested joins, yet are expected to deliver sub-second responses to those queries.

Database administrators are constantly creating and dropping indexes, changing data storage allocation parameters, defining and redefining partitions, and using various other techniques to boost performance. They must also monitor the database for any problems and perform preventative maintenance. In addition to this, the processes involved in applying software upgrades and patches can be highly disruptive and time consuming.

Although most modern RDBMSs, including Oracle Database, have developed many features and functions that are designed to lessen the burden of these activities, it remains the case that, at the highest level of size, complexity, and performance criticality, meeting the SLA for the database requires a set of skills generally regarded as something of a black art. And of course, even the best of DBAs can make mistakes – and those mistakes can cause poor performance and unscheduled downtime.

Approaches to Automating Database Management

The most common approach of major RDBMS vendors to these problems has involved optimizing the database internals, so that it simply runs faster while requiring less tuning. Approaches such as in-memory columnar structures have helped in this regard, but not entirely addressed the problem. They have worked well for the most common cases of databases, but for the aforementioned databases,

more is needed. Sometimes, an RDBMS will have software that analyzes its own structure and provides guidance regarding tuning (sometimes this capability is called "a DBA in a box"). These approaches are nearly always based on known issues and problem patterns, and the software tends to be driven by a static set of rules. Also, even RDBMSs that have been called "autonomic" usually require the DBA to actually take the tuning action because a measure of human judgement is generally required.

Machine Learning and Operations Automation

Databases and other types of enterprise software have had heuristics for years that provide various levels of operations automations. Oracle is no exception to this. What is new is the use of machine learning algorithms that replace the heuristics. There are numerous reasons for this – lack of sufficient amounts of data needed to train an ML model, lack of compute power to train the model effectively and in near real time, and lack of a sufficient variety of data coming from different types of users and use cases that helps to broaden the applicability of the algorithms.

Over the past three years, changes in data collection, compute power, and cloud-based services have mitigated the challenges outlined previously. We now have various IT services that are somewhat AI driven, using machine learning, natural language processing, and question-answer technologies to help build knowledge bases and predictive models that can be used to automate various parts of IT operations.

Oracle Autonomous Database: Automation with Machine Learning

What It Does

Oracle Autonomous Database is different from various "autonomic" RDBMSs in that it covers and automates the full range of database tuning and maintenance tasks. It can do this in part because it is delivered as a cloud service; the database server runs in a very well-defined and controlled environment, and the user has no hand whatsoever in the management of the database. Only when the management software has complete control of the database and its environment can it operate in this manner. In addition to providing all tuning and maintenance functions, most of which are automated, this service also provides regular software patching and upgrading, so the user is always running on the latest software, and knows that, for instance, the most recent security patches have been applied. Taken altogether, when one considers both the database tuning and maintenance, and also the software maintenance, Oracle Autonomous Database is likely to save the average organization hundreds or, in some cases, thousands of FTE hours per year for each major database.

Oracle could not offer this service without applying some advanced technology to solve the more nettlesome tuning and maintenance problems. Adjustments based on rules and patterns are not sufficient, and the human intervention required to apply such maintenance would make this service too expensive to represent a significant improvement over the prevailing do-it-yourself model. To scale this service, Oracle developed automated database tuning and maintenance based on machine learning.

How the Machine Learning Works

The core of any AI-enabled application is its learning algorithms. Developers and organizations are just beginning to understand and utilize the power of machine learning. Machine learning is defined as a toolbox of computer algorithms that create a model of some aspect of the world, enabling computers to learn, discover, predict, and improve automatically without explicit programming. Deep learning (DL)

is a type of supervised machine learning based on neural network algorithms, which has seen commercial success.

Although machine learning libraries have been around for decades and have been offered as part of many of the world's statistical packages, including IBM's SPSS, SAS, and so forth, the use of machine learning by enterprises hasn't been widespread until recently because these algorithms require a lot of data and a lot of compute power. However, many leading technology firms have been using machine learning tools over the past few years to improve programs in many areas, such as image recognition and programmatic advertising, as well as product and story recommendations and now IT and database operations.

Supervised machine learning begins with examples of training data paired with identifying labels (e.g., right or wrong, positive or negative) selected from the categories to be learned. Using these pairs of example data and labels ("training data"), the system learns parameters of statistical models that it can then generalize to unlabeled examples of data items that were not seen in the training data ("test data"). In most cases, the learned models improve over time via a feedback loop that adjusts the model parameters to better reflect additional sets of training or production data. The performance of a learned model can be measured by simple prediction accuracy or by the business metric the learned model is designed to support. Performance depends on the degree to which the training data matches the real world, the choice of algorithm, the algorithm's parameters, and the quantity of data.

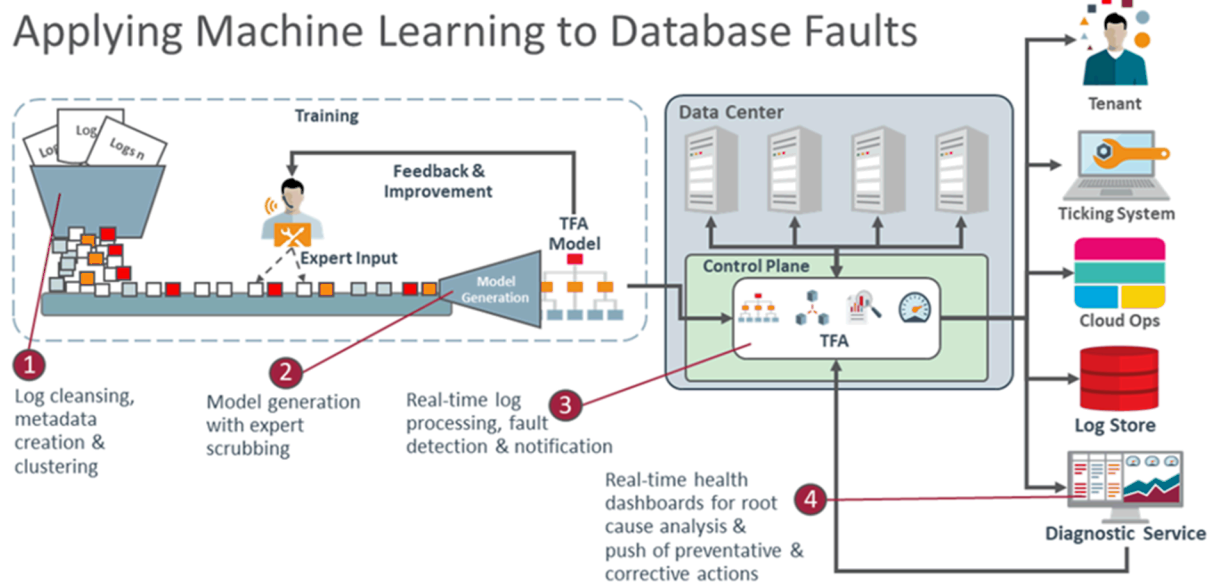
With these tools, the end game is all about making applications smarter by using special libraries containing self-learning algorithms, which, when unleashed on a data set, can learn and self-program to solve various types of problems. The combination of high-performance compute resources, tremendous amounts of data, and these ML/DL libraries is solving problems and challenges without the need to resort to programming. These ML/DL libraries and technologies are being used for an ever-wider array of use cases, from image recognition and disease diagnosis to pricing optimization and product recommendations.

As an example of how this works, Figure 2 shows how Oracle is going about the process of automating error handling within its Autonomous Database product. Operations log data is used as the source of information that is then used to train a machine learning model with some level of supervisory insight being provided by database experts. That model is then moved into production, helping to identify and handle errors, performing root cause analysis, and providing preventative and corrective recommendations for action to solve the various problems that emerge as part of the daily operations of the database.

While this is a relatively simple example, Oracle is working to automate most, if not all, aspects of database operations with its new Autonomous Database product – and automated error handling is just one example of where AI and machine learning are being applied in the new product offering.

FIGURE 2

Automated Error Handling



Source: Oracle, 2017

One Customer's Experience: A Search Service Provider

11880.com is one of the biggest directory assistance and directory listing companies in Germany. 11880.com is a small and medium-sized business (SMB) customer that was interested in becoming a beta tester for Oracle's Autonomous Database because of potential cost savings. We spoke with Christian Maar, the company's CEO and Vorstand member.

Maar indicated that 11880.com runs Oracle Database in support of a variety of both transactional and analytic applications. The transactional applications include the SAP application suite and Siebel CRM, and the analytics are driven by Oracle Business Intelligence (BI). "We were eager to be a beta tester for the [Oracle] Autonomous Database, because it will be possible to have heavy cost savings," he told us. "The database architects want to fully concentrate on architecting highly sophisticated databases. They want to design leading-edge technology and are happy that all the DBA stuff is put out of their hands."

Maar indicated that the company was an early beta tester, and the first two phases of the rollout were a bit "rough." It had been testing some of the workloads in comparison with a memory-optimized analytic RDBMS, and the other RDBMS was winning easily in the first two phases. But, in the third phase, not only did the Oracle Autonomous Database outperform the memory-optimized competitor but it did so while requiring no human effort at all.

Maar expects the benefits to 11880.com to come both in the form of process cost for running the databases and in smoother business continuity and better performance. At the time this document was written, 11880.com had not yet moved the Oracle Autonomous Database into production but the test results are very promising indeed.

ADVICE FOR THE TECHNOLOGY BUYER

Oracle Autonomous Database delivers major benefits to customers in two forms: labor savings by automating database tuning and problem detection and reduced downtime due to unplanned outages and software/system maintenance. It should provide Oracle Database users with a major incentive to move to the Oracle Cloud. Currently, the Oracle Autonomous Database addresses analytic workloads such as data warehouses – but we can expect a transaction-centric version later this year.

In evaluating this product, technology buyers should consider the following:

- This product is likely to evolve rapidly over the next couple of years. If it looks like a good fit for your organization, you should look at getting your feet wet now, but watch how it evolves before going all in.
- As more clients sign-on, Oracle Autonomous Database will learn from their cases, adding a richer range of situations and responses and making the system "smarter." So, expect results to improve over time.
- Competitors are likely to introduce their own autonomic features, so if your investment is in another vendor's product, check with that vendor to see what the plan is.
- A move to the cloud involves a radical shift in the IT topology and cost structure. Be sure you anticipate all the issues involved in moving the database and its applications.
- There are also personnel implications. Many DBAs with whom we have spoken in the past have complained about the tuning and maintenance tasks they have to perform because those tasks keep them from the stack of user requests for business-relevant changes to the database. Those DBAs will be happy with this technology. Others, however, regard their detailed knowledge of Oracle Database tuning to be their source of job security. They may not be so sanguine about adopting this technology.
- Communication with others is important. See what they are doing and what lessons there are to learn. This may affect your timing in moving to Oracle Autonomous Database.

LEARN MORE

Related Research

- *Cloud Database Adoption Trends* (IDC #US42074617, November 2017)
- *IDC FutureScape: Worldwide Analytics and Information Management 2018 Predictions* (IDC #US42619417, October 2017)
- *IDC Market Glance: Cognitive/AI Systems, Search, and Content Analytics, 3Q17* (IDC #US43098517, September 2017)
- *Market Analysis Perspective: Worldwide Cognitive/AI Software Platforms, 2017* (IDC #US42072717, September 2017)
- *IDC's Worldwide Semiannual Cognitive/Artificial Intelligence Systems Spending Guide Taxonomy, 2017: Update* (IDC #US43031517, September 2017)
- *Is Cognitive/AI Technology Adoption Real?* (IDC #DR2017_T5_DS, February 2017)
- *IDC PlanScape: Implementation of Cognitive Systems* (IDC #US41477516, June 2016)
- *Intelligent Personal Assistant Adoption for Work* (IDC #US41477816, June 2016)

- *IDC Innovators: Machine Learning-Based Text Analytics, 2016* (IDC #US41312116, May 2016)
- *IDC TechScape: Cognitive Systems Technologies, 2016* (IDC #US41005816, February 2016)

Synopsis

This IDC Perspective reviews Oracle's recent Autonomous Database announcement and examines how one customer is planning to make use of the automation capabilities provided in its day-to-day business processes.

"AI and machine learning are increasingly pervasive technologies that are being used now in a wide range of products and services," says Carl Olofson, research vice president, Data Management Software at IDC. "Oracle has recently launched its Autonomous Database product that uses AI and machine learning to automate operations and management of the product. Organizations should consider Oracle's Autonomous Database cloud offering if they are evaluating cloud database offerings and are looking for a fully managed database service from a proven database vendor. The future of autonomous products like Oracle's Autonomous Database is bright, but organizations need to also assess the impact these types of products will have on their IT staff."

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-community.com
www.idc.com

Copyright Notice

This IDC research document was published as part of an IDC continuous intelligence service, providing written research, analyst interactions, telebriefings, and conferences. Visit www.idc.com to learn more about IDC subscription and consulting services. To view a list of IDC offices worldwide, visit www.idc.com/offices. Please contact the IDC Hotline at 800.343.4952, ext. 7988 (or +1.508.988.7988) or sales@idc.com for information on applying the price of this document toward the purchase of an IDC service or for information on additional copies or web rights.

Copyright 2018 IDC. Reproduction is forbidden unless authorized. All rights reserved.

