

Oracle's second-generation cloud is designed to be enterprise grade

Publication Date: 27 Feb 2019 | Product code: INT003-000329

Roy Illsley



Ovum view

Summary

One of the key challenges for any organization looking to move mission-critical workloads to the cloud is that most of the hyper-scale cloud providers designed their clouds for scale and not specifically to securely run mission-critical workloads. The key capabilities required for mission-critical workloads are a high-level of security that protects the corporate data and a degree of automation that simplifies the complex management challenges. Oracle announced at Oracle OpenWorld 2018 a second-generation cloud architecture that is designed around the core principle of being secure and autonomous.

Oracle's second-generation cloud architecture separates the control plane from the user's data

The growth of cloud, according to Ovum ICT Enterprise Insights, is set to change in 2019/20 in terms of the workloads that organizations will be running in the cloud. This will increasingly be the mission-critical core systems such as ERP, CRM, and databases. Oracle Cloud is differentiating itself from other cloud providers such as AWS and Azure in terms of its cloud architecture's ability to meet the exacting requirements of running these mission-critical workloads. The second-generation cloud does not allow the Oracle control plane to coexist alongside the user's data and therefore creates an "air gap" between them.

Oracle announced it has released Oracle Cloud Infrastructure (OCI) as a second-generation cloud architecture that is needed as more mission-critical workloads move to the cloud. The challenge for current cloud architectures is that they were designed to be multitenant and as such were perfect for supporting multiple different customers and scaling as demand dictated. However, although this architecture may be ideal for enabling cloud providers to deliver services to customers cost-effectively, it does not address some serious security concerns. Oracle, with its next-generation cloud architecture, has separated the control plane from the customer's data. This separation of customer's and provider's data provides two key security benefits. First, Oracle cannot see customer data and has no way of accessing it without the customer's permission and help. This delivers on privacy and addresses many other concerns customers in other countries have about the US Patriot Act and the ability of US-based companies to access their data if required to by US agencies. Second, having the provider's control data on a separate set of servers, with controlled port access open for allowed communications in one direction only, means any malicious attacker who gets access to a customer's data cannot also reach other customers or the provider's control plane. This level of security is precisely what is needed as the cloud begins to be used for more mission-critical and core business applications, which Ovum's data suggests will be in 2019/20.

New cloud offerings are designed to expand the use cases at an improved price point

Oracle announced that it would be offering a set of AMD-based compute instances based on the AMD EPYCTM processor. Oracle becomes one of the first public cloud providers to have a bare metal version with AMD EPYCTM processors. It is expected that these instances will cost \$0.03 per core

hour. This price represents a significant saving compared to the Intel-based instances that Oracle and other cloud providers currently offer. One of the criticisms of AMD processors has been their performance in terms of memory bandwidth, but the AMD EPYC can deliver >269Gbps. This makes it a very good general-purpose processor that meets the needs of big data analytics – a memory bandwidth-intensive operation.

In addition to the bare metal offering, the AMD offering also features the one, two, four, and eight core VM shapes. These instances are generally available in Oracle Cloud Infrastructure's US East Ashburn region today and will be available in London by the end of October and in other US and European regions by the end of the year.

The other two key announcements at Oracle OpenWorld were the introduction of a high-performance computing (HPC) shape with instances powered by Nvidia's Tesla V100 GPUs (which were released at Nvidia's GPU Tech Conference in March) and the introduction of clustered networking. Using Oracle's remote direct memory access (RDMA) protocol, the clustered networking will support networks of up to 100Gbps, providing an ultra-low-latency and high-bandwidth solution that is ideal for HPC, databases, big data, and AI workloads.

The autonomous database as a service will enable organizations to exploit greater business value and service availability

The core of any organization's value is its corporate data, and the majority of this data is held in databases. However, these data repositories are complex and expensive assets to manage while ensuring data is both secure and available. As organizations continue their journey to cloud computing, moving these corporate repositories is now beginning to tax the minds of CIOs. Currently, the cloud providers enable organizations to migrate these databases from on premises to a cloud model, but these cloud database services are still rudimentary in terms of the service capabilities they offer. Oracle has an extensive heritage in the database market, and while it has been accused of being slow to move this expertise from an on-premises model to a cloud model, its reasons are understandable: Oracle did not want to reduce the quality of service customers could expect from a database as it moved to the cloud.

The Oracle Autonomous Database Cloud (OADC) is a new concept that Oracle is rolling out throughout 2018/19. The OAC is built using AI and machine learning and is designed to improve availability while reducing the cost of managing these complex assets. Initially, the OADC will be provided for data warehouse deployments, but it will be expanded to cover OLTP, NoSQL, and graph use cases. The core premise of OADC is to transform the way databases are managed when they move to cloud to improve some challenging operational activities such as tuning databases to ensure optimum performance, patching and updating databases to ensure security and compliance, and auto-scaling databases to meet business demand. The cost saving associated with automating these activities, typically performed by DBAs, will allow these expensive resources to be used for more value-adding data science activities, enabling organizations to exploit their data assets. The other main advantage of this autonomous approach is it will provide a guarantee of service availability, with Oracle quoting no more than 30 minutes of downtime a year; this figure includes any time required for scheduled maintenance and updates.

Appendix

Author

Roy Illsley, Distinguished Analyst, Infrastructure Solutions

roy.illsley@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

The contents of this product are protected by international copyright laws, database rights and other intellectual property rights. The owner of these rights is Informa Telecoms and Media Limited, our affiliates or other third party licensors. All product and company names and logos contained within or appearing on this product are the trademarks, service marks or trading names of their respective owners, including Informa Telecoms and Media Limited. This product may not be copied, reproduced, distributed or transmitted in any form or by any means without the prior permission of Informa Telecoms and Media Limited.

Whilst reasonable efforts have been made to ensure that the information and content of this product was correct as at the date of first publication, neither Informa Telecoms and Media Limited nor any person engaged or employed by Informa Telecoms and Media Limited accepts any liability for any errors, omissions or other inaccuracies. Readers should independently verify any facts and figures as no liability can be accepted in this regard – readers assume full responsibility and risk accordingly for their use of such information and content.

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

ovum.informa.com

askananalyst@ovum.com

INTERNATIONAL OFFICES

Beijing

Boston

Chicago

Dubai

Hong Kong

Hyderabad

Johannesburg

London

Melbourne

New York

Paris

San Francisco

Sao Paulo

Shanghai

Singapore

Sydney

Tokyo

