Technology Oracle

Sponsored by Oracle

How Melbourne University combines clouds to compute

Peter Shadbolt

Feb 15, 2022 – 8.00am

The next decade stands to see incredible change due to the multiplier effect of various emerging technologies coming together, including AI, ML, IoT and cloud.

On the one hand, IoT has been promised to bring incredible change to both industries and everyday lives, given the millions of connected devices in society.

Yet, the volume and capacity of data produced has been congesting networks and overpowering the systems of companies and researchers.

The era of multi-cloud architecture. **Getty Images**

On top of this, the multiplier effect on cloud has paved the way into the era of multi-cloud.

“The cloud has changed massively since its first inception more than a decade ago, becoming more comprehensive, but at the same time more flexible.” says Simon Kaye, Director, Cloud Transformation and Innovation at Oracle.

Today many enterprises now are boosting their bottom line by engaging a multi-cloud architecture that distributes their workload across multiple platforms, often reducing costs at the same time as mitigating risks associated with a single provider.

Rajkumar Buyya, Redmond Barry Distinguished Professor and Director of the Cloud Computing and Distributed Systems (CLOUDS), at the University of Melbourne. **Supplied**

At the heart of this trend is the understanding that no two cloud services, as with no two businesses, will ever be exactly alike.

Where previously enterprises were locked into cloud services that offered different policies, pricing and sometimes an opaque view of their full functionality, today there’s more choice.

What this means is that organisations are now in a better position to mix and match, effectively molding cloud services to fit their specific needs - swapping out the functions that are not a good fit for their enterprises while retaining those that do.

“When first conceived, cloud computing was regarded a dark art – people looked at it and said ‘wow it’s like the mainframe’. But what they didn’t realise at the outset was the investment that was put into delivering the cloud both economically and how they could use this investment to extract and deliver new capabilities,” he says.

The time when people simply shifted their data centre to the cloud and on to someone else’s balance sheet has changed too, he says.

“Now it’s evolved to where companies and organisations are thinking, ‘How do I build a custom cloud solution and/or multi to get the best of all worlds.’ It’s really

about how you simplify and enrich your work practices, but you use someone else's investment to deliver that."

Multi-cloud, he says, is increasingly an established part of the business landscape, delivering speed and performance - which suits complex database applications that are not well served by traditional cloud infrastructure providers.

It also brings scale because deploying more than one cloud increases a company's ability to scale up as needed and move data around as required; and it lowers prices (using multiple clouds gives businesses leverage in negotiating both price and contract) but competition also drives investment and innovation.

Where Oracle fits in, says Kaye, is as the central hub for all of the business's cloud and data requirements.

"CIOs are asking for three things: How do I take my existing investments in Oracle platforms and technologies and achieve more with them? How do I aggregate capacity at the lowest possible cost? How do I integrate all this with other systems that I'm availing myself of now and into the future?

"You might have a SaaS (software as a service) interest that are not Oracle's. You might have other cloud interests that are not Oracle's. How do you make all of that actually work?

"Oracle has for decades been associated with systems of record and companies and organisations are building systems of engagement and insight that rely on this data out to various different cloud solutions."

Rajkumar Buyya, Redmond Barry Distinguished Professor and Director of the Cloud Computing and Distributed Systems (CLOUDS), at the University of Melbourne has engaged Oracle Cloud Infrastructure, as part of the widely available and secure multi-cloud service that underpins its FogBus2 development and research activities.

The curiously named project is part of the latest in "edge" and "fog" computing - new computing paradigms that provide resource-limited internet of things (IoT) devices with scalable computing and storage resources.

Compared to cloud computing, edge/fog servers have fewer resources, but can be accessed with higher bandwidth and less communication latency. This becomes especially powerful in production control or operation technology type environments.

The aim, ultimately, is for researchers in academia and industry to utilise research and optimise IoT applications, devices and particularly the data flow which threatens to overwhelm current network and computing resources.

The university chose to partner with Oracle because of the IT giant's support for ARM processors – similar to those found in mobile phones, and a perfect fit for IoT data processing ensuring predictable network performance.

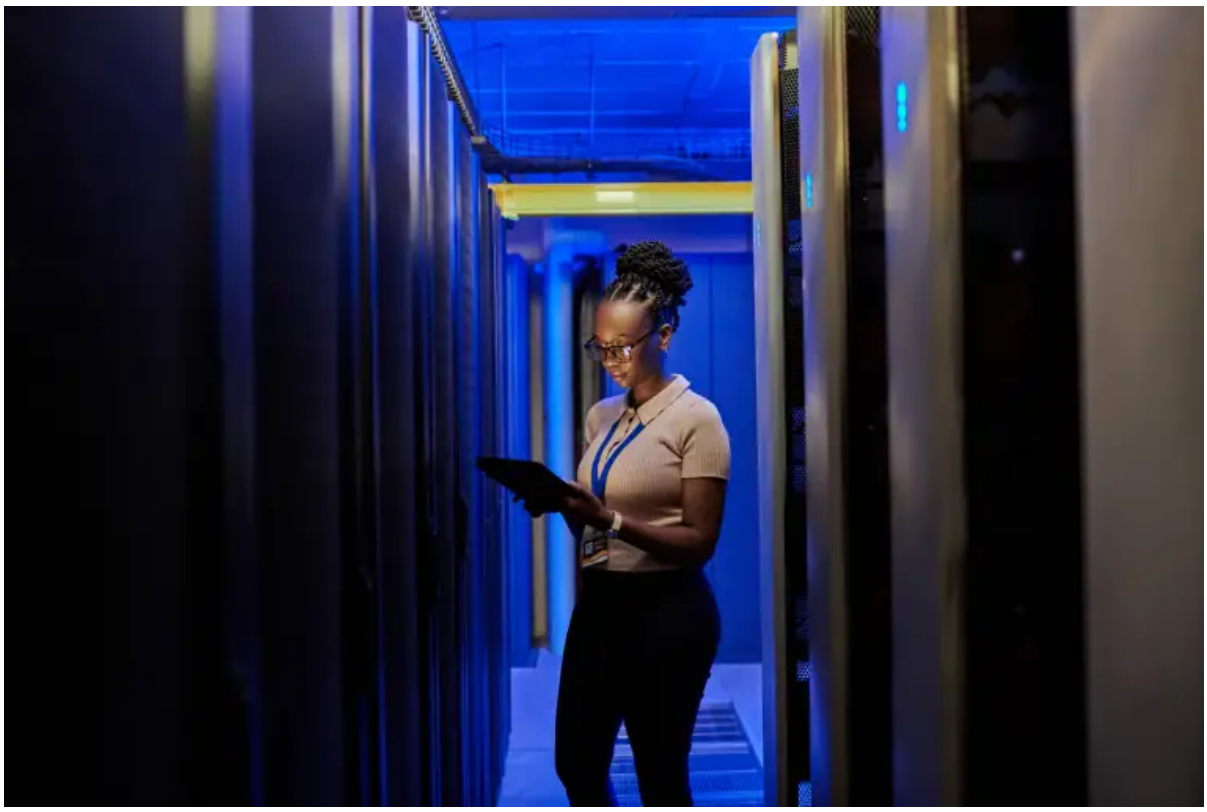

Multi-cloud delivers speed and performance. **Getty Images**

Oracle Cloud Infrastructure connects the edge processing capabilities, currently using a mixture of an open-source database and Oracle Autonomous Data Warehouse, with the core processing which also utilises Oracle's converged database and Oracle Machine Learning.

This uses pre-built algorithms to allocate networking and compute resources to processes, and helps reduce the need for expensive experts.

"This multi-cloud solution helps ensure we have the right location for data processing," Buyya explains.

"For instance, if you were developing a smart transport system then sensors are placed to capture the road conditions.

"When data is captured it needs to be analysed in a timely manner. If you send the data to the core there will be more latency in getting your answer, so the algorithm will determine the time sensitivity of that data, and the need to use the edge cloud computing environment so that you can get your answer quickly."

Alternatively, if you are taking in images from an MRI, there could be thousands and thousands of data points.

"Sending all that information to the network might congest the system," he says.

"So the purpose of FogBus2 is to look at the use of algorithms to analyse frames at the edge, so that only frames that are sequentially different from each other are transmitted to the cloud for further processing, so as to enable the network choose the closest or the right cloud for your analytics. It's this sort of thing that multi-cloud brings to the table," he says.

Additionally, the autonomous database allows the university to remove a lot of the need for database administrators as it self-patches, self-secures and self-repairs.