

Oracle Cloud Infrastructure Blog

Technical Solutions

University of Melbourne deploys Arm-based, distributed container framework for fog computing on Oracle Cloud Infrastructure

May 18, 2022 | 9 minute read



Kellsey Ruppel

Principal Product Marketing Director, OCI



Karthik Sukumar

Principal Cloud Specialist Engineer at OCI



The authors would like to thank Dr. Rajkumar Buyya, Redmond Barry Distinguished Professor and the director of the university's Cloud Computing and Distributed Systems Lab, and Mohammad Goudarzi, graduate researcher at the University of Melbourne's CLOUDS Lab, in the School of Computing and Information Systems at the University of Melbourne for their contributions.

The





world's compute and networking infrastructures face a growing challenge: floods of data originating from billions of Internet of Things (IoT) devices coming online. The sources of that data include medical equipment; video cameras that require image recognition; sensors throughout factories, airplanes, smart vehicles, and tractors; and countless consumer devices, from refrigerators to light bulbs. The volume of data captured from connected devices threatens to overwhelm networks and compute resources, making it impossible to extract useful insights.

The [University of Melbourne](#) in Australia prides itself in bringing together some of the best and brightest students and researchers from around the world to tackle complex problems such as this one in innovative ways.

Goals for cloud migration

Computer science researchers at the University of Melbourne are tackling these problems by building FogBus2, which is an open source, container-based, distributed framework for collecting data from cameras, ECG devices, laptops, smartphones, or most any other IoT devices. When fast results are needed—for example, acting on readings from a medical sensor—edge servers near the IoT device that run Raspberry Pi and Jetson processors analyze that data quickly. When more compute is needed but response time is less important, the data goes to the cloud for analytics. The goal is to manage traffic and provide compute infrastructure that can help make sense of such data.

The researchers needed to build FogBus2 on a platform that's multicloud, widely available, secure, and works in real-world implementations, not just a controlled lab environment.

“With all these billions and billions of IoT devices pushing data to the cloud, the latency would be very high,” says Rajkumar Buyya, Redmond Barry Distinguished Professor, and the director of the university’s Cloud Computing and Distributed Systems Lab. “That is where we start moving toward this model of fog computing, where we can harness the resources at the edge of the network and the cloud. As a research laboratory, it is exciting to show that our effort is not just functional in our own controlled lab infrastructure, but that it works in the real world. Working with Oracle helps us demonstrate that capability.”

Suite of Oracle products used

[OCI](#) includes all the services needed to migrate, build, and run IT in the cloud—from existing enterprise workloads to new cloud native applications and data platforms. The researchers used the following OCI services and technologies:

[Regions](#): An OCI region is a localized geographic area that contains one or more data centers, called

availability domains. Regions are independent of other regions, and vast distances can separate them (across countries or even continents).

Identity and Access Management (IAM): IAM enables you to control who can access your resources in OCI and the operations that they can perform on those resources.

Virtual cloud networks (VCNs) and subnets: A VCN is a customizable, software-defined network that you set up in an OCI region. Like traditional data center networks, VCNs give you complete control over your network environment. A VCN can have multiple non-overlapping CIDR blocks that you can change after you create the VCN. You can segment a VCN into subnets, which can be scoped to a region or to an availability domain. Each subnet consists of a contiguous range of addresses that don't overlap with the other subnets in the VCN. You can change the size of a subnet after creation. A subnet can be public or private.

Security lists: For each subnet, you can create security rules that specify the source, destination, and type of traffic that must be allowed in and out of the subnet.

Route tables: Virtual route tables contain rules to route traffic from subnets to destinations outside a VCN, typically through gateways.

Virtual machines (VMs): The Compute service enables you to provision and manage compute hosts in the cloud. You can launch compute instances with shapes that meet your resource requirements for CPU, memory, network bandwidth, and storage. After creating a compute instance, you can access it securely, restart it, attach and detach volumes, and terminate it when you no longer need it. Ampere Altra processors come with 80 cores per CPU, with all cores capable of running at the maximum frequency of 3.0 GHz consistently. Each core comes with its own 64-KB L1 I-cache, 64-KB L1 D-cache, and 1-MB L2 D-cache.

Bastion: OCI Bastion provides restricted and time-limited secure access to resources that don't have public endpoints and that require strict resource access controls, such as bare metal and virtual machines, Oracle MySQL Database Service, Autonomous Transaction Processing, Oracle Container Engine for Kubernetes, and any other resource that allows SSH access. With the Bastion service, you can enable access to private hosts without deploying and maintaining a jump host. In addition, you gain improved security posture with identity-based permissions and a centralized, audited, and time-bound SSH session.

Autonomous Data Warehouse: Oracle Autonomous Data Warehouse is a self-driving, self-securing, self-repairing database service that is optimized for data warehousing workloads. You don't need to configure or manage any hardware, or install any software. OCI creates, backs up, patches, upgrades, and tunes the database.

Oracle Machine Learning: Oracle Machine Learning (ML) is an in-database service that helps improve data scientist and developer productivity, and reduce their learning curve by using familiar open source-based Apache Zeppelin notebook technology. Notebooks support SQL, PL/SQL, Python, and markdown interpreters for Oracle Autonomous Database, so users can work with their language of choice when developing models. Python users accelerate model building and execution by using Autonomous Database's built-in Python environment as a high-performance computing platform. Built-in Auto ML provides automated algorithm and feature selection, as well as model tuning and selection.

Migration path and Oracle solution

The University of Melbourne moved its FogBus2 research platform to OCI, allowing computer science researchers throughout Australia to capture, analyze, and make predictions from IoT data.

Using a three-tier architecture—consisting of IoT devices, a cluster of edge devices featuring Raspberry Pi and Jetson Nano, and a multicloud network—the university’s researchers can instantly ingest, analyze, and make inferences about IoT data coming from healthcare monitoring systems, surveillance cameras, mobile devices, and more.

Learn more about the University of Melbourne's journey to Oracle Cloud:

Within its VCN on OCI, the University of Melbourne has a public subnet, in which researchers can process workloads on four VMs. Two VMs use Arm processors for

University of Melbourne Analyzes IoT Data from the Edge t...

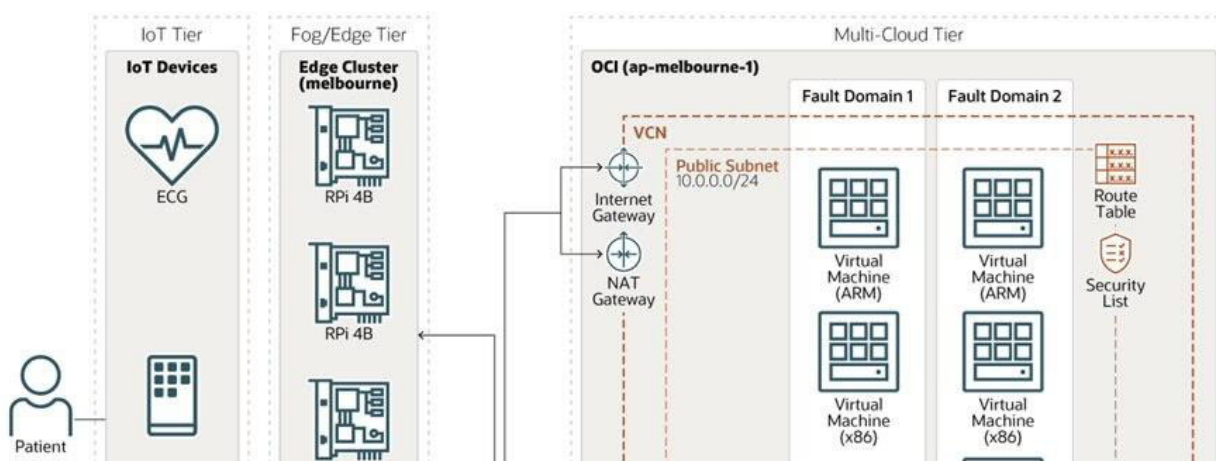


computationally intensive and latency-sensitive workloads, such as those involving video analytics or electrocardiogram readings. Another two VMs use Intel x86 CPUs for lighter, more structured workloads, such as patient health records. The workloads are deployed across two fault domains for high availability. Each fault domain includes an Arm-based VM and an x86 VM. To secure access, a bastion host is deployed in the public subnet, providing restricted and time-limited access to the university's VMs.

Researchers chose to run FoguBus2 on OCI, including Oracle Autonomous Data Warehouse and Oracle Machine Learning.

The University of Melbourne runs its backend systems, including an Autonomous Data Warehouse and a VM running an Intel x86 CPU, in a private subnet. To maintain a schedule for performing system tasks, the University of Melbourne research team uses Autonomous Data Warehouse to log the performance of the edge devices, CPU cores, and memory capacity. The team then uses in-database Oracle Machine Learning to run algorithms that can predict the optimal runtime for those tasks.

The following diagram illustrates this architecture.



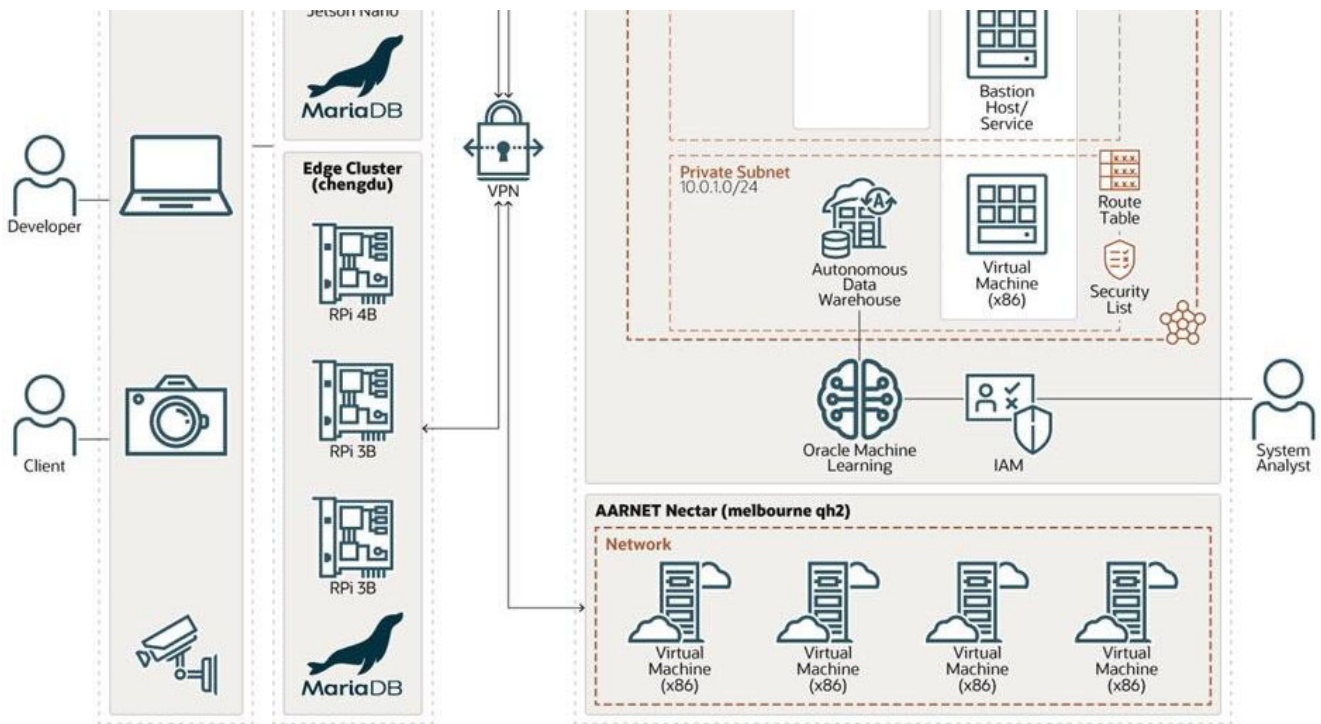


Fig 1: University of Melbourne's Reference Architecture

The university is also part of the [Oracle Arm Accelerator](#) program, which provides a one-year free trial to select university researchers, open source developers, industry partners, and private companies to run workloads on OCI [Ampere A1 Compute](#) and similar services.

The Arm processor offers support for a VM's hundreds of cores, which provides an opportunity to massively scale operations within each VM. Autonomous Data Warehouse supports built-in artificial intelligence, containerized software, microservices, and edge computing.

The researchers also tapped [Oracle Cloud Free Tier](#), which provides free services that any developer can use to experiment and find the best architecture to support applications, before deploying at a large scale.

Putting IoT to work: FogBus2 real-world example

Video image recognition, such as assessing surveillance cameras for criminal activity, is the kind of application that FogBus2 is built to facilitate.

“If you look at a citywide deployment, with cameras in multiple places, there could be thousands and thousands of data points, and if each computation takes just two minutes to analyze, that can take several hours on a single virtual machine,” says Mohammad Goudarzi, graduate researcher at the University of Melbourne's CLOUDS Lab in the School of Computing and Information Systems.

A typical video camera captures information at 30 or 60 frames per second. Sending all that information to the network might congest the system. So instead, FogBus2 uses algorithms to analyze frames on the edge servers, and only frames that are sequentially different from each other are transmitted to the cloud for more processing.

For example, if a store manager wants to assess COVID-19 mask compliance by shoppers, they can set up cameras to capture images of people coming into a shopping center and then analyze faces to see

which shoppers are not wearing masks. An application like that needs to be fast, with results available in real time and analyzed on the network edge, Goudarzi says. Store managers could then direct staff to ask those shoppers to wear masks.

Analysis of protective face covering requirements is more compute-intensive and can be processed at the edge or in the cloud.

“Our framework can support applications that are computationally intensive or latency sensitive. That is one of its selling points,” Goudarzi says.

Results achieved with OCI

Oracle Cloud Infrastructure helps FogBus2 deliver fast results. The Arm-based OCI services provide FogBus2 with rapid service startup time, 15% faster than services running on other architectures that developers tried. Arm also delivers 15% to 20% performance increases and response times in applications, including video recognition and other latency-sensitive applications that require high parallelization.

Learn how the University of Melbourne uses Oracle Autonomous Data Warehouse and Oracle Machine Learning:

Autonomous Data Warehouse delivers the scalable CPU and storage needed for IoT workloads, which generate sudden bursts of traffic. The data warehouse’s automation increased developer



productivity, allowing the university to develop FogBus2 in one year.

Autonomous Data Warehouse’s native support for ML reduces the need for developers to build and manage ML capabilities in applications. Oracle Machine Learning allows researchers to quickly manage data visualization within a database, with only a few clicks and lines of code and without the need to build separate backend services.

Oracle Machine Learning allocates networking and compute resources for data analytics within the FogBus2 framework. Autonomous Database compiles network information, such as latency, data rate between devices, availability of hardware resources such as CPU and RAM utilization, and software information such as response time and which containers are running on which devices. Data is stored in flat tables, with each record just a simple row of data, to allow FogBus2’s remote logger container to append and read data quickly.

FogBus2 implements a three-tier framework:

The IoT tier at the edge comprises the actual IoT devices gathering data, such as cameras and medical equipment.

The second tier is edge clusters, running Raspberry Pi and NVIDIA Jetson Nano processors, with a MariaDB database. The edge clusters manage local analytics for applications requiring low latencies.

Data from the edge clusters flow over a VPN to a third tier, which is running in the OCI region in Melbourne, with Arm and x86 processors running Autonomous Data Warehouse and Oracle Machine Learning.

In summary, the University is experiencing the following results with OCI:

Lightweight and distributed container-based framework for integrating heterogeneous IoT-enabled systems with FogBus2 and cloud servers

Fast and low-overhead deployments of applications using containerization and near linear scaling for CPU-bound workloads

Scheduling, scalability, resource discovery, and dynamic profiling mechanisms that assist IoT developers to define and deploy their targeted IoT applications

Integrated scheduling, scalability, and resource discovery policies

No constraints on communication topology between its entities and support for different topologies such as mesh, peer-to-peer, and client-server

Next steps

In
the



future, the university plans to integrate container-orchestration techniques to automate management. They also plan to adopt privacy-preservation support for user and FogBus2 server private information. The university will do this by using private peering options through multiple private connections between the Australia Academic and Research Network (AARNet) and OCI data centers without using a VPN or Network Address Translation (NAT). Finally, they plan to integrate ML techniques to analyze

the current state of the FogBus2 computing environment.

Explore more Oracle Cloud Infrastructure customer [technical case studies](#) and try [Oracle Cloud Free Tier](#).



Kellsey Ruppel

Principal Product Marketing Director, OCI

I'm a Principal Product Marketing Director in the Oracle Cloud Infrastructure GTM team, where I'm responsible for OCI customer references, messaging, positioning, field enablement, business development, and marketing strategy and execution across OCI.

[Show more](#)



Karthik Sukumar

Principal Cloud Specialist Engineer at OCI

[← Previous Post](#)

Resources for

- About
- Careers
- Developers

Why Oracle

- Analyst Reports
- Best CRM
- Cloud Economics

Learn

- What is Customer Service?
- What is ERP?

What's New

- Try Oracle Cloud Free Tier
- Oracle Sustainability

Contact Us

- US Sales 1.800.633.0738
- How can we help?
- Subscribe to Oracle Content

Investors	Corporate Responsibility	What is Marketing Automation?	Oracle COVID-19 Response	Try Oracle Cloud Free Tier
Partners	Diversity and Inclusion	What is Procurement?	Oracle and SailGP	Events
Startups	Security Practices	What is Talent Management?	Oracle and Premier League	News
		What is VM?	Oracle and Red Bull Racing Honda	