

# ARCOS establishment roadmap

A nationally coordinated approach to supporting containers, Kubernetes, and multi-cloud interoperability

Prepared by the ARCOS establishment steering group: Steve Quenette (Monash University), Paul Coddington (ARDC), Mark Gray (Pawsey Supercomputing Centre), Steven Manos (BioCommons), Ryan Fraser (AARNET), ARCOS Working Group.

V1.1 2020-07-01

## Introduction

Kubernetes is the emerging open source standard for deploying and managing scalable and portable container-based cloud-native applications. It has become the most commonly-used approach for multi-cloud application deployment, which for digital research infrastructure enables communities to develop tools & commons that can be deployed where the community member has container friendly resources. Being at the juncture of infrastructure (cloud/hpc) and applications (platforms), implementing Kubernetes is non-trivial, but made substantially easier when Kubernetes is consistently operated for each cloud and/or application and best-practice approaches are used.

The ARDC's Storage and Compute theme's primary role is to deliver underpinning computing capabilities to the ARDC competencies, and the need to better support containers and Kubernetes had been identified in the theme roadmap. The [proposal for a "Kubernetes Core Services"](#) - to establish and operate for 3 years a scalable multi-cloud platform and expertise to support federated, distributed cloud-native research applications - was submitted to the EOI phase of the 2019 Platforms Open Call. This was a signal that there was strong community support for this service, and ARDC's Compute and Storage theme accordingly reprioritised plans to support this as an ARDC activity. The ARDC's Compute and Storage theme and the national community have collaborated to develop the EOI into a "Kubernetes Core Services establishment roadmap" for ARDC's consideration.

## Purpose

This proposal seeks to emulate the transformational impact for organisations and research through almost 10 years of the ARDC (OpenStack) Core Services competency, by establishing a Kubernetes Core Services. This will be a national team of experts in Kubernetes. This includes supporting related open source software for cloud native computing, engaging with the international Kubernetes and Cloud Native Computing communities to ensure Kubernetes is developed to meet Australia's

needs, establish an innovation and operational best practices amongst the national stakeholders, and co-operate on developing part of the national technical eResearch ecosystem.

A Kubernetes Core Service will transform the deployment and operation of computing platforms and tools, especially those being developed through the ARDC Platforms program on the Nectar Research Cloud. Moreover, it is an evolution of ARDC's Storage and Compute theme, where in addition to seeding capacity, it enables Australian researchers to become increasingly hard-infrastructure agnostic. Researchers, institutions and domains can focus their investments on tooling (platforms and data commons), knowing that the provision of compute and storage capacity can be determined as and when needed. This approach will underpin and amplify the continued Australian participation and leadership in international research tools and commons.

## Consultation

At eResearch Australasia 2019 (October) an establishment steering group was formed, with representatives from Monash University (Steve Quenette, the EOI lead), ARDC (Ian Duncan and Paul Coddington), Pawsey (Mark Gray), and the BioCommons (Steven Manos).

A community call to action was organised with a [workshop](#) at the Pawsey Supercomputing Centre on 11 March 2020. There were 30 attendees, with representatives from eResearch infrastructure providers including ARDC, Pawsey, NCI, AARNet, CSIRO, Nectar Research Cloud Nodes (Monash University, Intersect, TPAC, and QCIF), universities, NCRIS capabilities (TERN, Auscope, and BPA) and ARDC Platforms (EcoCommons, Australian Characterisation Commons, Australia's Scalable Drone Cloud, and the BioCommons BYO Data Platform). Participants presented use cases of the application of Kubernetes (see appendix), and several distinct findings emerged.

### Consultation findings:

- Containers are now prolific. They are primarily used to simplify packaging and deployment of tools and platforms;
- Container orchestration and interoperability across multiple compute platforms are still challenging problems for the developers and operators of research tools / platforms / data commons;
- Kubernetes is the dominant platform for container orchestration and multi-cloud deployment;
- Kubernetes is hard to use and has a steep learning curve;
- Researchers and developers becoming Kubernetes experts is an unnecessary burden, reducing their efficiency at their core competency. Yet without Kubernetes, those tools will have limited use and exposure;
- No one researcher, developer group or organisation has a vested interest to support Kubernetes beyond their immediate need. Supporting national coordination and standardisation of underpinning infrastructure, and in this case hard-infrastructure agnostic infrastructure, is a natural role for the ARDC.
- Make it easier for Australian developers and operators of research platforms to deploy on multiple compute platforms including the Nectar Research Cloud.

# Proposal

There was strong support from all attendees to establish the next Core Service capability, in partnership with the ARDC:

- As a national activity that provides expertise, standards, interoperability, and best practices for containers, Kubernetes, and multi-cloud interoperability;
- As a collaboration between multiple organisations;
- Engaged with and contributing to the relevant international communities and open source software projects;
- a project charter and project plan should be developed and endorsed by an inclusive working group, including those from the EOI, call to action, and beyond;
- a small steering committee, including those from the establishment group, should oversee the development of a project charter and projects plan, and the resources committed to the project;
- Concerted coordination and drive is necessary, requiring dedicated staff, and ARDC is best placed to provide this resource.

The project charter and project plan are being developed. The community is seeking feedback from the ARDC as to its involvement.

## Recommendation

This activity is strongly aligned with the ARDC Storage and Compute theme strategy. This work will make it easier for users of the Nectar Research Cloud, including several ARDC Platforms projects, to deploy and operate applications and platforms, and to ensure they work interoperably across different Nodes of the Research Cloud, other sector-based cloud infrastructure (e.g. at NCI and Pawsey), and commercial clouds. Several key users of the Nectar Research Cloud including ALA, TERN, Auscope, Biocommons, and Ecocloud have indicated that they want to run, or are already running, services on both the Nectar Research Cloud and other cloud infrastructure including commercial cloud.

We propose that ARDC support this initiative, and this support takes the form of:

- Commit to a 3 year initiative in line with the Research Cloud Refresh and Platform activities (maintain certainty from the perspective of the user)
- Provide the project coordinator role (approx 0.5 FTE)
  - possibly within the ARDC Outreach and Engagements team
- Establish a Kubernetes technical lead (1 FTE), to maintain ARDC as a leader in the Australian community
  - preferably within the ARDC Core Services team
- Recognise some effort will be required from other members of the Core Services team
- Participate in the charter's objectives to develop a national roadmap of requirements and service provision
- Contribute funding for supporting roles at ARDC and/or the initiative partners (with requirements for co-investment), as determined collaboratively with the ARDC through the roadmap, that by principle fill a gap that the community would not otherwise provide, but drive uptake of peak and systemic usage and strategic impact. (1-5 FTE, years 2 & 3)

Some of the current members of the ARDC Core Services team have some limited expertise in Kubernetes which they are growing, but for a technical lead for this project we would expect to recruit a Kubernetes expert.

# Appendix 1

## Kubernetes Use Cases and Requirements

Summary of use cases and requirements presented at the Kubernetes workshop at the Pawsey Supercomputing Centre, 2020-03-11.

- **BioCommons**

- Biocommons Pathfinder work package on software and containers has been exploring requirements for deployment of life sciences software and how this is being done in the international bio community, e.g. the ELIXIR project
- Workshop held at Pawsey on 2020-03-11 discussed requirements and planned approach, which is to follow the approach used by ELIXIR and Galaxy of using biotools, bioconda, biocontainers to package all software as containers
- Pawsey are using Singularity containers to deploy software on cloud and HPC
- We should support deployment and use of platforms and tools across lots of different compute platforms - best approach for this at the moment looks like K8s
- We should be making this easy for people
  - Standardise on what are we providing, when do we upgrade versions, etc
  - K8s is hard to use, Docker Swarm is easier but it's going away
- Don't want researchers to be devops experts, we need to help them

- **Ecocloud**

- Ecocloud runs on openstack, use k8s for deployment
- duplicated ecocloud to Tinker (HASS DeVL)
- had problems with limitations in Research cloud. E.g. neutron networking not fully deployed, node instabilities, performance with storage backends, missing shared file systems.
- K8s deployment with adapted kubespray to cater for Research Clouds specifics.
- led to ARDC Platforms EOI proposal for Launchpad
- was thinking about maybe ARDC just sorting out k8s clusters at multiple cloud nodes/sites, 1 FTE for Ecocloud just sorting out deployment

- **Characterisation and Imaging Platforms**

- [Presentation slides](#)
- hard to get software installed on HPC, want to make this easier
- so far looked at containers but not orchestration - apart from slurm
- different people have different approaches to building containers, need best practice
- many users created docker containers that are not portable
- have made best practice info for creating containers openly available

- **Drones Platform**

- Building a cloud-native infrastructure for drone data platform
- K8 as a core infrastructure to run a scalable data processing applications, web-based analysis and visualisation applications
- Platform run across multiple ARDC regions including at QCIF, Monash, CSIRO, etc

- want people to be able run data processing pipeline as containers where data is stored
- Offer container-based data processing that can run at different regions including commercial clouds
- **Kubernetes infrastructure at Monash**
  - making software available as docker containers
  - need QA for processes to create containers - so they don't get hacked e.g. for bitcoin mining
  - best practice around containers is important - want to trust researchers
  - different users having different approaches to K8s - e.g. using Magnum, or not
  - aiming to run Jupyter notebooks in k8s clusters, so can use smaller VMs
  - can binpack things better if use containers
  - Want to use containers and K8s for XNAT and Mytardis as well
  - CogStack (EHR -> NLP Engine) pipelines on k8s
  - Run machine learning pipelines on k8s e.g. kubeflow pipelines to process data, train/retrain with different models and hyper-parameters, test and display results
  - Deploy/serve trained ML models
  - Run/provide services for teams e.g. jenkins, authentication, monitoring
- **Auscope**
  - Interest in containers and K8s from the simulation side of Auscope
  - Australian research codes that are used globally - dockerising these, for teaching as well
  - Need to address issues around provenance of code, citations, etc - aiming to get academic credit
  - similar investigation in biocommons discussion
  - Zenodo registry for linking publications, data, and code/containers/workflows
  - using binder and mybinder as well
  - have been using AWS
  - With AVRE they have been following what Louis is doing - he's leading the pack
- **CSIRO Mineral Resources (Online Processing Toolkit)**
  - creating a platform to consolidate development
  - Researcher code is containerised and made available for execution via the platform.
  - aiming to be more portable
  - Using gitlab CI as a workflow engine, taking advantage of gitlab k8's deployments and auto scaling runners.
  - key features around auto scaling, using serverless workloads
  - using minio not specific buckets like gcp or s3, can connect to cloudstor
  - object storage universally used but different interfaces in different clouds so this affects portability
  - Early stage of development
- **AARNet**
  - everything containerised
  - using Rancher
  - S3 built on K8s - wrote own helm charts
  - aiming to move Cloudstor to K8s

- Swan is run on k8s at AARNET, every new jupyter instance is a k8s pod
  - Swan is an owncloud service developed by CERN, but CERN don't use k8s for deployment
  - Align AARNet's cloud offering to be consistent in so it can be used by many
  - Provide vanilla compute with K8s
  - Need QA processes for "shared" or created containers to ensure security
- **TERN**
    - Run all applications and data processing pipelines as containers
    - Use K8s as core platform to run client-facing applications and services containers
    - Applications interface with ARDC Object storage, Volume storage, CloudStor and RDS storage;
    - Run K8s nodes in multiple ARDC regions (QRISCloud and TPAC)
    - Use k8s to run automated Apache Airflow tasks for data onboarding to the TERN data infrastructure as and when data are available;
    - Applications running on K8s interface with stand-alone applications running in their own clusters including Geoserver, Postgres, Apache Spark.

Jericho

Been doing manual deployments of K8, for Jupyter workloads. Doing a cookoff of differ deployment tools at the moment.

**TERN**

Running 2x k8 clusters on the research cloud there, across multiple nodes. Issue with neutron stability. Moved to K8 to get this layer of abstraction in place. Eventually move, or run hybrid.

Why move? Well there's data sitting in AWS S3 - range of reasons why Public cloud would be of interest.

Monash - running in to problems with advanced networking. MyTardis running on K8 - slowly porting services to this. Running 3 xNat instances - working on scripts to move into K8 to help with scalability.

AIS - setting up federate xnat repos.

Pawsey - Can't scale to the userbase we expect use existing direct access interfaces. Core interest - make our job as operators easier. Need to come up with a operating model and a constrainable support load, that is ultimately somewhat scalable. Future of being responsible for more stuff they are currently responsible for.

AARNet - on prem - Perth, Melbourne, Brisbane and Canberra. 1 main cluster and a couple of dev clusters. Using Rancher (open source).

K8 - enabler of multisite. Standard methods and workflows for creating containers and building containers.

Where do containers come from? User born, or 'golden containers'. Trusted containers.

## **Shared challenges discussion - Friday 22 May 2020**

K8 - architecture that platforms are adopting -

NIF - want to have a trusted repository for containers. Paula - work for NIF, but also involve with MASSIVE and Microscopy Australia. It's the characterisation community - containers and software tools for image analysis in AU. Differences in clusters, so containers are not transferable. Have a day to tag containers that are available - there's a limit to the information that's publicly available.

ARDC - next call - architectural guidance could be provided by the ARDC- Could K8 be an architectural approach.

What problem are we trying to solve? K8 is hard- want to make it easier to support. Come up with a model we can support, a model that can provide a better experience for researchers working with research workloads.

Desire to avoid variant re-invention. Let's make it easy for people to do stuff, like setting up a registry. A challenge to resolve -

Not a lot of exploitation at the moment on what K8 offers. Right now, used as a dependency approach and as a deployment/sharing approach. Cross infra use case as well - across provider types.

First steps towards an early win:

How do we get stuff in to a container

How do we share it around

Locating/tagging/metadata

Advantage of using a domain (say CVL) repo - more trust than say a public repo

Trust - operators

CVL - they trust each other, they trust their own repo. What we are trying to sold her- multiple groups trusting each other's stuff (outside of their community). -

Establish some guidelines - if you have a container - who is the author, when was it updated, etc. - metadata - extra resources to help trust that container. BioCommons has the same challenge here.

Automated testing - have a gate for different clusters - passes a gate to run on a different platforms. That just a technical solution to running (or not).

Nectar golden images - well it's core services that produces.

But there's the issue of human trust.

NCI convo - need to build it bit by bit



Need a framework for trust? CICD framework for people to plug in to?  
What information do people need to know - it's going to differ.

Kieran - Mark - Paula - Manos - develop the summary.

Paula - Share responsibility of maintenance - Users need trust immediately. HPC admins won't risk installing from an unknown source.

## Appendix 2

### List of Kubernetes workshop attendees

Crystal	Chua	AARNet
Ryan	Fraser	AARNet
Paul	Coddington	ARDC
Andrew	Treloar	ARDC
Sam	Morrison	ARDC
Jake	Yip	ARDC
Steven	Manos	Biocommons
Brian	Davis	Biocommons/NCI
Samuel	Bradley	CSIRO
Jens	Klump	CSIRO
Sven	Dowideit	CSIRO
Sarah	Richmond	Griffith/QCIF
Israel	Casas	Intersect
Toan	Nguyen	Monash
Swe	Aung	Monash
Jerico	Revote	Monash
Lance	Wilson	Monash
Steven	Quenette	Monash
Andrew	Howard	NCI
Allan	Williams	NCI
Jenni	Harrison	Pawsey
Marco	De La Pierre	Pawsey
Audrey	Stott	Pawsey
Sarah	Beecroft	Pawsey
Mark	Gray	Pawsey
Wei	Fang	UNSW
Michael	Mallon	UQ/QCIF
Siddeswara	Guru	UQ/TERN