A reservation system in the Nectar Research Cloud for GPU and large memory instances

20 October 2022
14.05 - 14.25

PRESENTED BY
Paul Coddington

Australian Research Data Commons

NEW SERVICE - **why?**

- Growing demand for high-end compute infrastructure in the research sector
  - Seen in ARDC Platforms
  - Seen across institutions with demand for GPUs for ML, image processing, simulation
  - Requirement for large memory machines to handle big data sets and large scale analysis and simulation
- Not limited to one type of research discipline

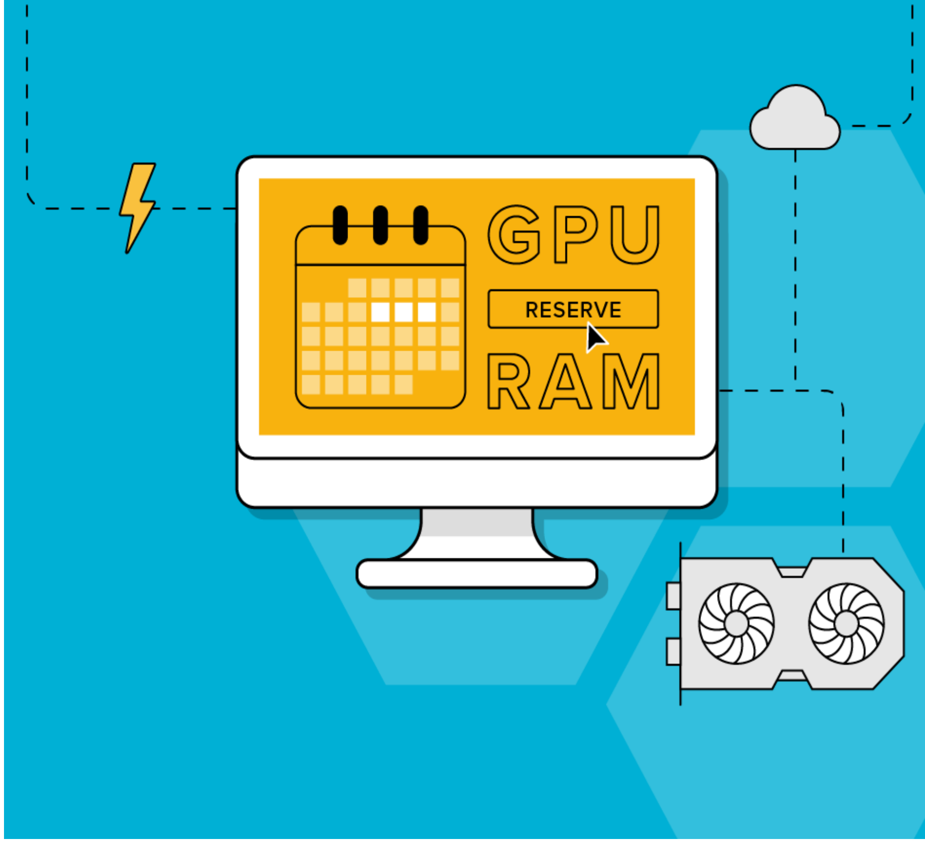

NEW SERVICE - **why?**

- GPU and large memory servers in the Nectar Research Cloud were
  - dedicated to Platforms projects or
  - for local Node use
- Want to provide a national IaaS for projects meeting national merit criteria
- Nectar infrastructure is provided at no cost to researchers - but GPU and large memory servers are very expensive
- Need to ensure high utilisation of these expensive and limited resources

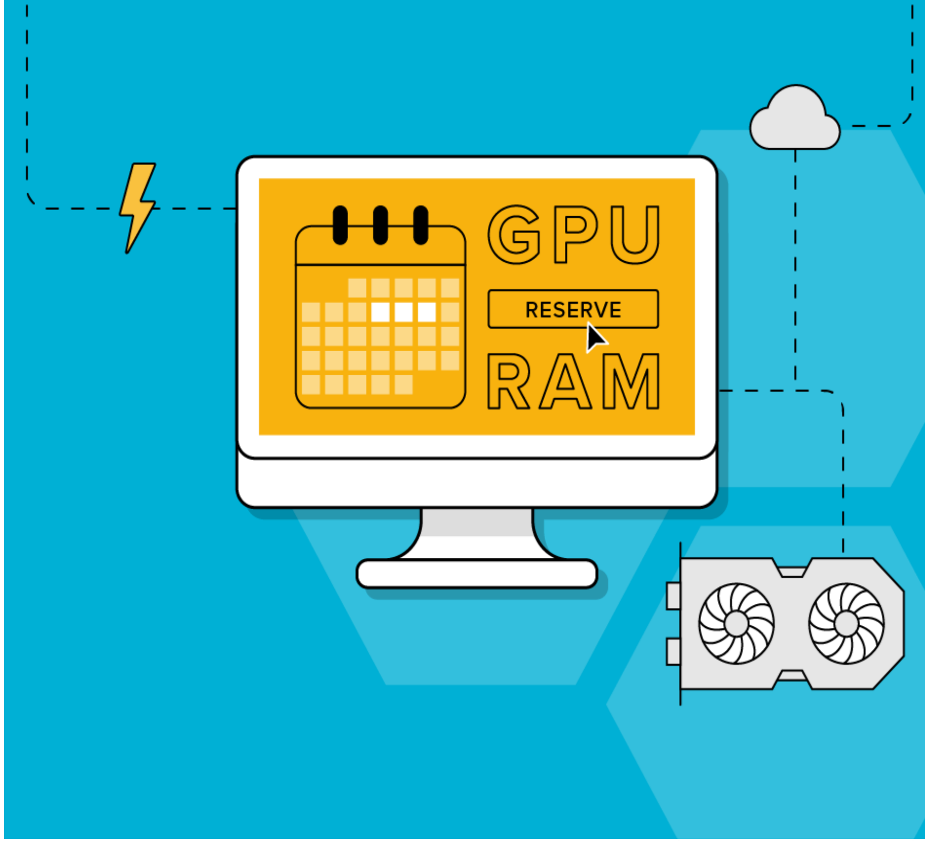

- **Solution - virtualisation and a reservation system**

Steps required to move to new service

1. Move from limit-based allocation quotas to usage based quotas
2. Design and develop a reservation system
3. Define and develop standard flavors for GPU and large memory virtual machines
4. Provide new GPU and large memory hardware at Nodes to underpin the service
5. Design and test with a Pilot Phase
6. Launch service

1. Move to Usage Based Allocations - Service Units

OVERVIEW

- ➢ Move from maximum capacity limit to credit/budget of service units (SUs) for the period of the project allocation
- ➢ This will then account for actual usage of resources where each resource has a specific cost
- ➢ Service developed for generic cloud allocations (all users)

WHY?

- *Maximum capacity allocation limits flexibility*
- *Adapt to bursty data analysis workflows*
- *Service designed for diversity of compute workflows*
- *More efficient use of available capacity*

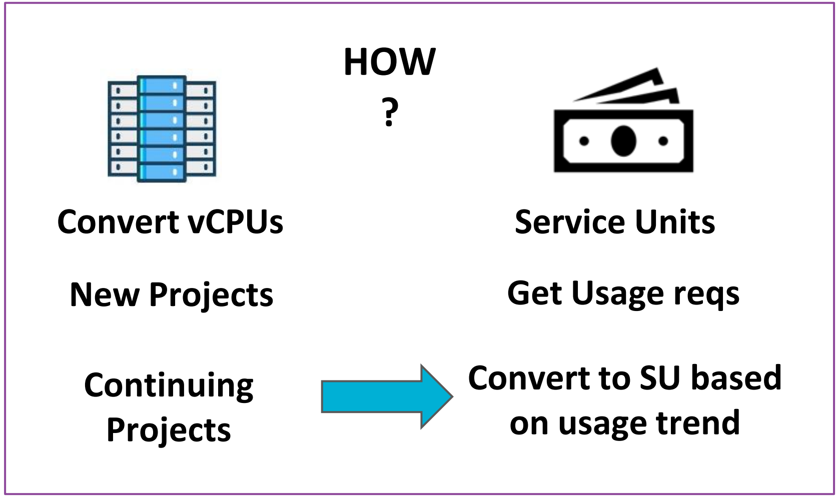

HOW ?

Convert vCPUs

Service Units

New Projects

Get Usage reqs

Continuing Projects

Convert to SU based on usage trend

Change to Dashboard User Interface (UI)

Allocation Usage

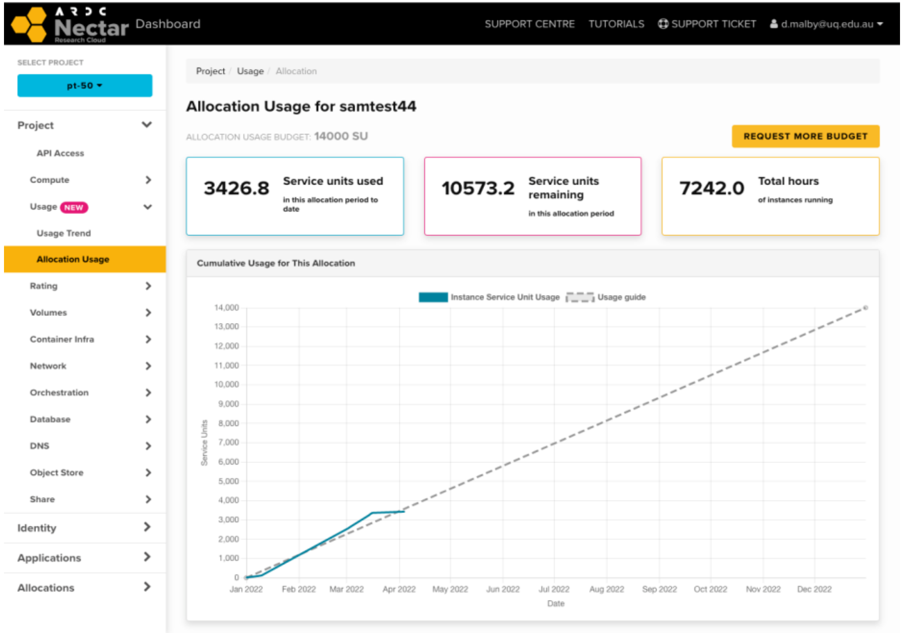

Allocation Usage for samtest44

ALLOCATION USAGE BUDGET: 14000 SU

3426.8 Service units used in this allocation period to date

10573.2 Service units remaining in this allocation period

7242.0 Total hours of instances running

Cumulative Usage for This Allocation

Usage Trend

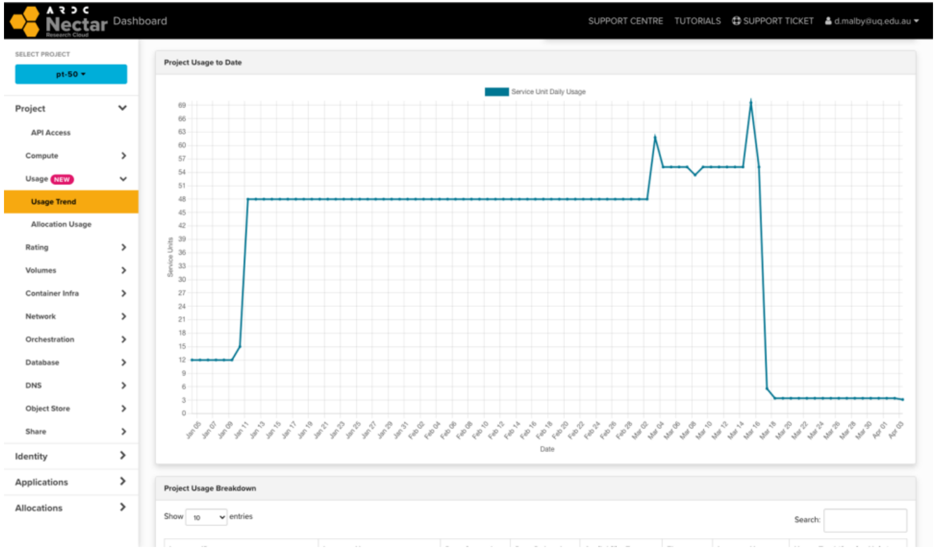

Project Usage to Date

Project Usage Breakdown

Show 10 entries

Search:

2. Design & Implement a Reservation System

Built on the OpenStack Blazar reservation service.

Users can reserve access to specialised high-end computing power in the Dashboard.

* Allocations must first be approved for reservations.

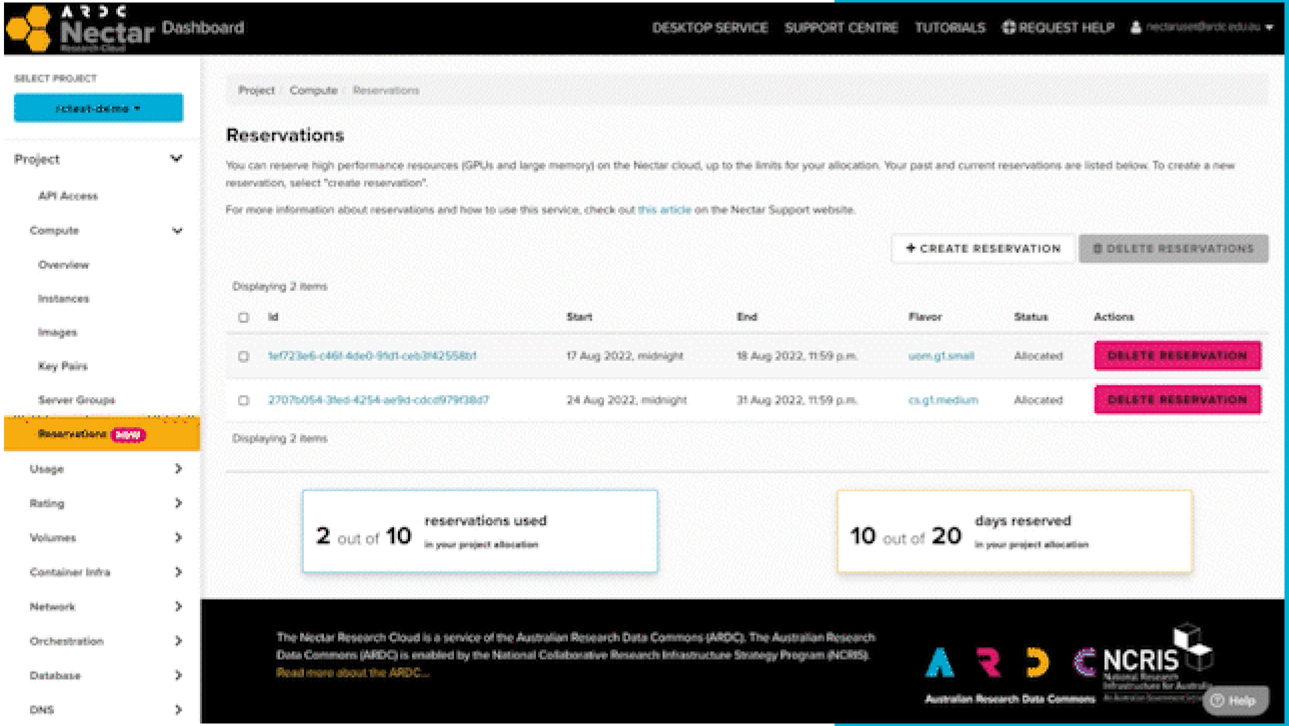

Reservations

You can reserve high performance resources (GPUs and large memory) on the Nectar cloud, up to the limits for your allocation. Your past and current reservations are listed below. To create a new reservation, select "create reservation".

For more information about reservations and how to use this service, check out this article on the Nectar Support website.

Displaying 2 items

□	Id	Start	End	Flavor	Status	Actions
□	1ef723e6-c46f-4de0-9fd1-ceb3f42558b1	17 Aug 2022, midnight	18 Aug 2022, 11:59 p.m.	uom.g1.small	Allocated	DELETE RESERVATION
□	2707b054-3fed-4254-ae9d-cdcd979f38d7	24 Aug 2022, midnight	31 Aug 2022, 11:59 p.m.	cs.g1.medium	Allocated	DELETE RESERVATION

Displaying 2 items

2 out of 10 reservations used in your project allocation

10 out of 20 days reserved in your project allocation

The Nectar Research Cloud is a service of the Australian Research Data Commons (ARDC). The Australian Research Data Commons (ARDC) is enabled by the National Collaborative Research Infrastructure Strategy Program (NCRIS). Read more about the ARDC...

SERVICE UNITS - in the Reservation System

☑ Enough SU budget is available.

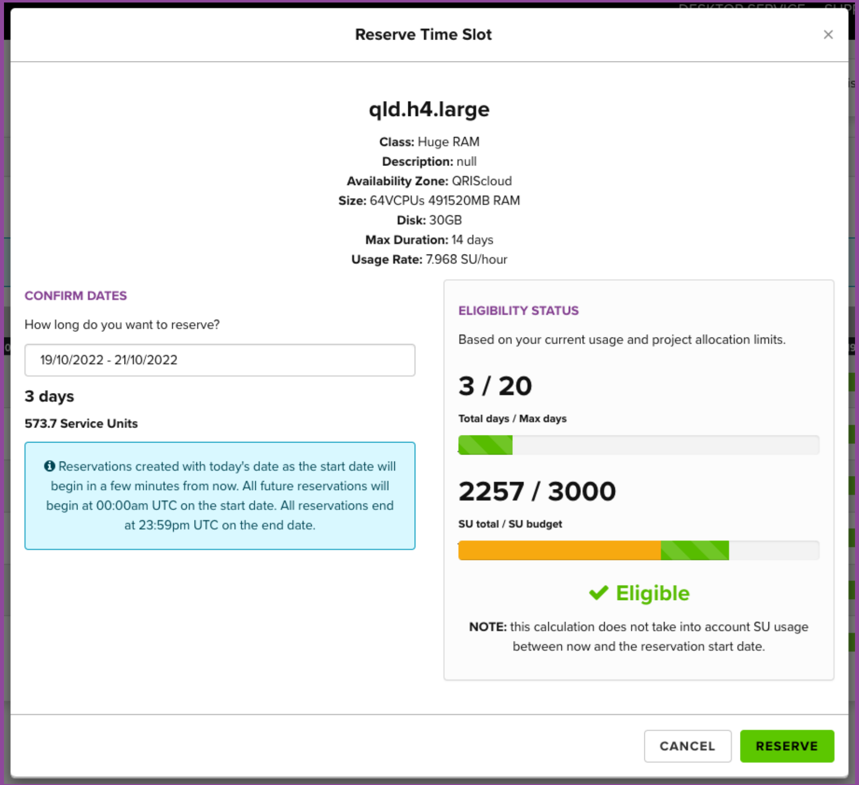

Reserve Time Slot

qld.h4.large

Class: Huge RAM
Description: null
Availability Zone: QRIScloud
Size: 64VCPUs 491520MB RAM
Disk: 30GB
Max Duration: 14 days
Usage Rate: 7.968 SU/hour

CONFIRM DATES

How long do you want to reserve?

19/10/2022 - 21/10/2022

3 days

573.7 Service Units

Reservations created with today's date as the start date will begin in a few minutes from now. All future reservations will begin at 00:00am UTC on the start date. All reservations end at 23:59pm UTC on the end date.

ELIGIBILITY STATUS

Based on your current usage and project allocation limits.

3 / 20

Total days / Max days

2257 / 3000

SU total / SU budget

✓ Eligible

NOTE: this calculation does not take into account SU usage between now and the reservation start date.

☒ Not enough SU budget is available.

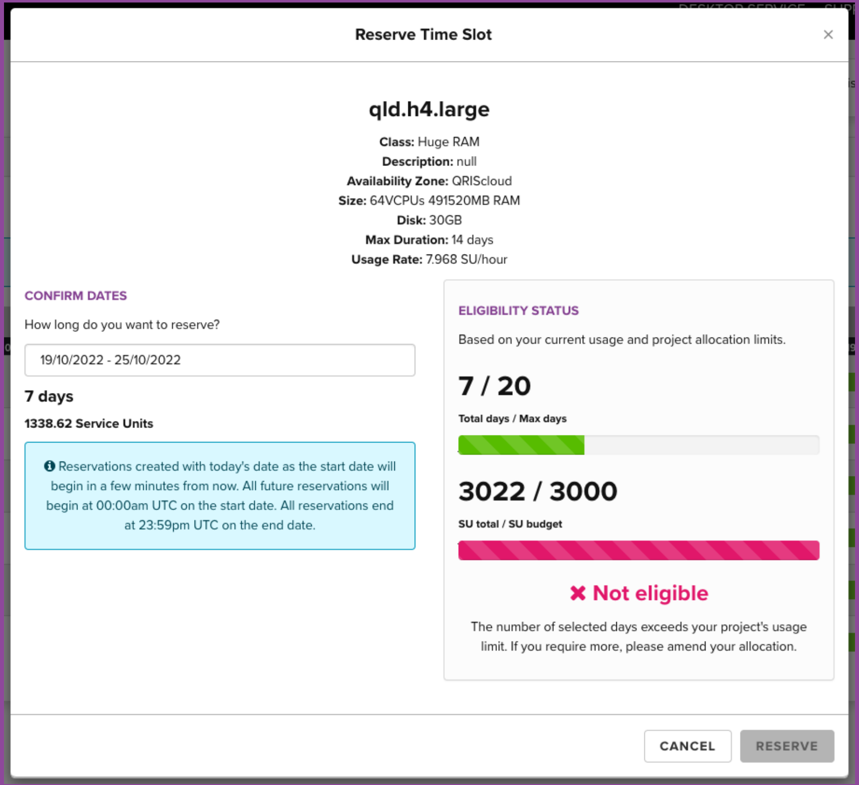

Reserve Time Slot

qld.h4.large

Class: Huge RAM
Description: null
Availability Zone: QRIScloud
Size: 64VCPUs 491520MB RAM
Disk: 30GB
Max Duration: 14 days
Usage Rate: 7.968 SU/hour

CONFIRM DATES

How long do you want to reserve?

19/10/2022 - 25/10/2022

7 days

1338.62 Service Units

Reservations created with today's date as the start date will begin in a few minutes from now. All future reservations will begin at 00:00am UTC on the start date. All reservations end at 23:59pm UTC on the end date.

ELIGIBILITY STATUS

Based on your current usage and project allocation limits.

7 / 20

Total days / Max days

3022 / 3000

SU total / SU budget

✖ Not eligible

The number of selected days exceeds your project's usage limit. If you require more, please amend your allocation.

3. Defining Flavors

The following standard flavor classes are offered on the Nectar Research Cloud:

- Tiny (t3)
- Balanced (m3)
- RAM Optimised (r3)
- CPU Optimised (c3)
- Preemptible (p3)
- **GPU Visualisation (g1) - A40**
- **GPU Compute (g2) - A100-80 sliced up 1/2 to 1/10**
- **Huge RAM (h4) - up to 128 vcpus and 960GB RAM**
- Huge RAM (h3)

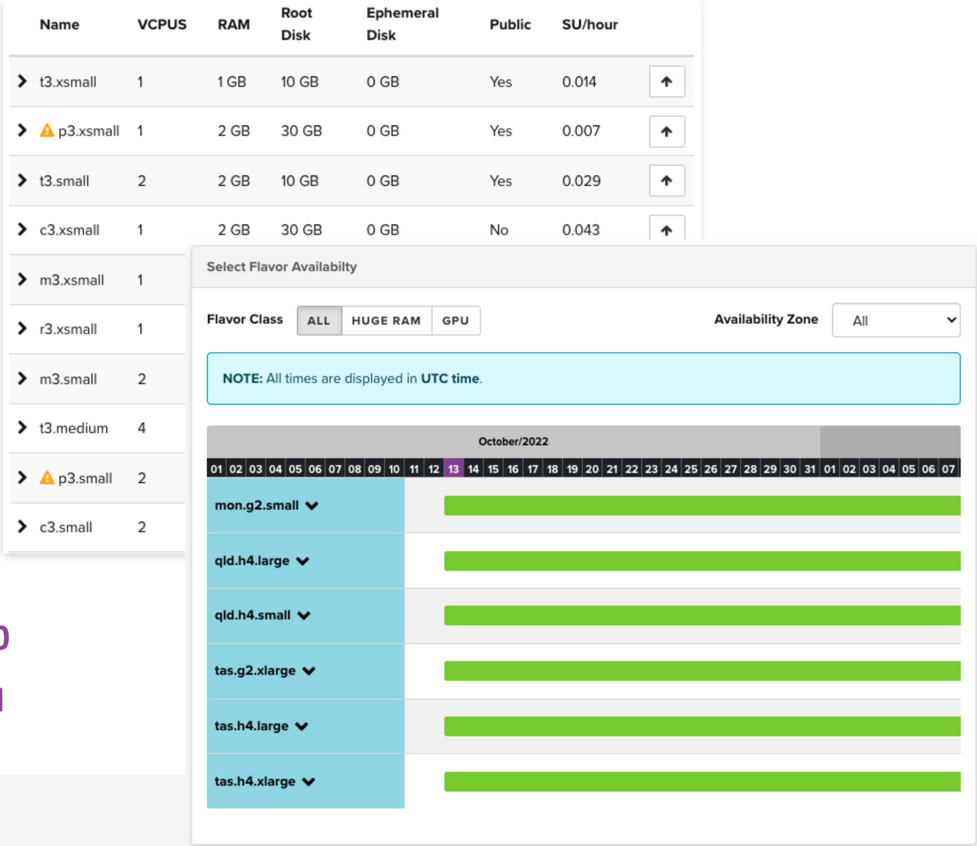

Name	VCPUS	RAM	Root Disk	Ephemeral Disk	Public	SU/hour
t3.xsmall	1	1 GB	10 GB	0 GB	Yes	0.014
⚠ p3.xsmall	1	2 GB	30 GB	0 GB	Yes	0.007
t3.small	2	2 GB	10 GB	0 GB	Yes	0.029
c3.xsmall	1	2 GB	30 GB	0 GB	No	0.043
m3.xsmall	1					
r3.xsmall	1					
m3.small	2					
t3.medium	4					
⚠ p3.small	2					
c3.small	2					

Select Flavor Availabilty

Flavor Class: ALL | HUGE RAM | GPU

Availability Zone: All

NOTE: All times are displayed in **UTC time**.

October/2022

01 02 03 04 05 06 07 08 09 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 30 31 01 02 03 04 05 06 07

mon.g2.small

qld.h4.large

qld.h4.small

tas.g2.xlarge

tas.h4.large

tas.h4.xlarge

Detailed list and recommended uses of Flavors available here:
https://support.ehelp.edu.au/support/solutions/articles/6000205341-nectar-flavors

4. Provision the hardware and licenses

- RFP to Nectar nodes asking for proposals for hardware based on their researcher requirements
- ARDC provided capex, Nodes provide opex as co-investment
- 16 GPU servers and 7 large memory servers from ARDC investment
- Additional servers from Node investment for use by Node members
- GPUs are a mix of
  - A100-80 GPUs (mainly for compute)
  - A40 GPUs (mainly for image processing and visualisation)
- All have large NVMe drives for fast local disk storage
- A pool of Nvidia licenses for virtualisation - VCS for A100 and VWS for A40

5. Pilot and Testing Snapshot

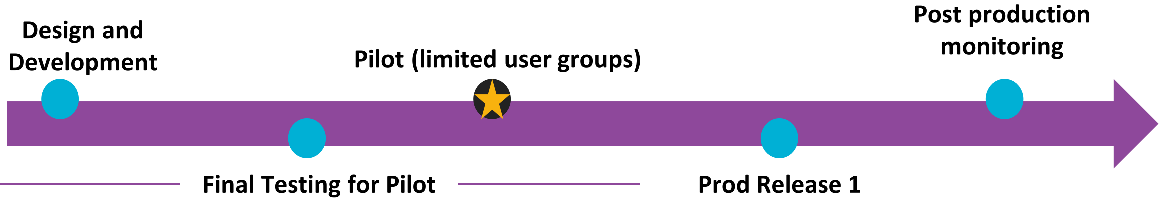

- Pilot environment for user testing released on 2nd Aug 2022
- 24 users across ~14 projects
- 2 GPU and 4 large RAM servers running with different Flavors
- 44 GPUs in total to be deployed
- ~248 vGPUs available for researchers

NATIONAL GPU SERVICE: **Benefits Realised**

Resource Allocation

- Fair allocation of resources that are limited & expensive
- Ensures that the limited resources are reserved and released when required
- Reserved access for training courses

Virtualisation

- Virtualisation enables improved GPU utilisation and more users
- First virtualised GPUs service on a national scale

Unique flavors

- A variety of flavors designed for research
- Provides access to large GPU flavours not yet available on any cloud platform in Australia
- GPU servers can be reconfigured to adapt to usage trends for different sized flavors

By end of 2023

248

vGPUs

Participating Nodes

University of Tasmania
Hobart

Monash University
Melbourne

QCIF
Brisbane

Intersect
Sydney

Swinburne University of Technology
Melbourne

6. Launch and expand

- Production environment released as ‘BETA’ on 13th September 2022
- Uplift in capacity as more infrastructure comes on line
- Available to all national merit research projects
- Capacity will be added as infrastructure comes on line
- Review and revise the mix of flavors, limits, etc

Learn more!
Webinar October 25th
2-3pm AEST

ACKNOWLEDGEMENTS

- Shubhra Dargar - project management
- Sam Morrison - reservation system
- Darcelle Maltby - web dashboard interface
- Sengor Kusturica, Rocky Yan, Andy Botting, Dylan McCulloch - GPU virtualisation, licensing, standard flavors
- Jo Morris and Sonia Ramza - user guides, user support, promotion
- ARDC comms team - communications and promotion
- And technical assistance from many Node operations staff

HOW CAN THE **ARDC ACCELERATE YOUR RESEARCH?**

Visit us at

eResearch Australasia - Stand 14

ARDC
Australian Research Data Commons

NCRIS
National Research Infrastructure for Australia
An Australian Government Initiative

ARDC is enabled by NCRIS

2022 Data Driven Research Impact

Download at **ARDC.EDU.AU ▶**

Subscribe to the **ARDC CONNECT** newsletter

THANK YOU

- ardc.edu.au
- contact@ardc.edu.au
- +61 3 9902 0585
- @ARDC_AU
- Australian-Research-Data-Commons

About the ARDC