



Research compute options at UoA

Martin Feller
Centre for eResearch



Before we start...

- If you have questions, please unmute yourself or use the chat. Please speak up in case I miss the chat
- Otherwise, please stay muted
- This talk is a general introduction to research compute options at UoA.
- Lunch and listen sessions...

ResBaz 2020 : Pick n Mix

23-27 Nov. Free, open, online - researchers & research postgrads
35 sessions, incl. Keystories - join us to listen over lunch



Mon. 12-1pm	<i>Welcome to ResBaz 2020 : Pick n Mix</i> https://auckland.zoom.us/j/99720152410 Passcode: 640143
Tues. 12-1pm	<i>Harnessing the disruptive nature of portable sequencing for community empowerment</i> https://vuw.zoom.us/j/95143657235 Passcode: 718144
Wed. 12-1pm	<i>From Classics to Computer Science and back again</i> https://vuw.zoom.us/j/98866515882
Thur. 12-1pm	<i>Performance - Music performance and project showcase</i> https://vuw.zoom.us/j/98265221971
Fri. 12-1pm	<i>Exploring the cultural horizons of open science: your research and your life</i> https://auckland.zoom.us/j/98447731180 Passcode: 404029

This talk

- Why or when would I need this?
- New Zealand eScience Infrastructure (NeSI)
- Virtual machines
 - What is a virtualisation?
 - When to use a virtual machine?
 - Options at University of Auckland
 - Nectar
 - Research Virtual Machine Farm
- Alternatives

Why or when would I need this?

- Offload computational tasks from laptop/desktop computer
- Simulations/Analyses run for days
- Laptop/desktop computer isn't good enough (CPU, RAM, storage)
- Use your own dedicated work environment
- Effective collaboration, e.g. shared development workspaces
- Reproducible research

New Zealand eScience Infrastructure (NeSI) 1/4



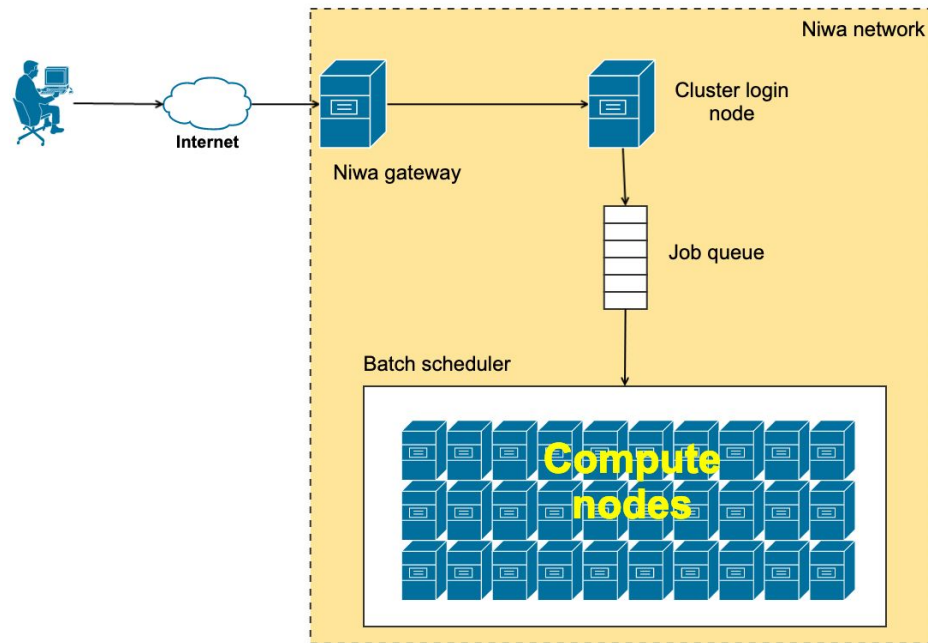
- National initiative, sponsored by MBIE, providing currently mostly shared high-performance computing (HPC) service (aka cluster computing)
- NeSI's collaborators invest in the infrastructure, and employ NeSI team members at their institutions



- HPC systems located in Wellington

New Zealand eScience Infrastructure (NeSI) 2/4

- Linux-based
- Batch-scheduling: you have to estimate runtime and resources for your analysis beforehand
- Access to these systems mostly on the command-line



<<Demo>>

New Zealand eScience Infrastructure (NeSI) 3/4

Pros and Cons

- + Your code can leverage many CPU cores across multiple computers
- + Multiple runs of the same code (parameter sweeps)
- No MS Windows workflows
- Technically challenging if you're not familiar with Linux
- Small-scale compute (not worth the effort)

New Zealand eScience Infrastructure (NeSI) 4/4

Getting started


- Request an account
- Request an allocation / project
- Documentation start: <https://support.nesi.org.nz/hc/en-gb>
- Weekly introductory workshop:
<https://support.nesi.org.nz/hc/en-gb/articles/360000042867>
6-Can-I-attend-an-Introductory-Workshop-

Virtual Machines (VMs)

Virtualisation

*“A **virtual machine** is a program on a computer that works like it is a separate computer inside the main computer”*
(Wikipedia)

- We can simulate a whole computer including the operating system
- Several virtual computers can run on one larger, more powerful computer



Virtual vs. Real Computer

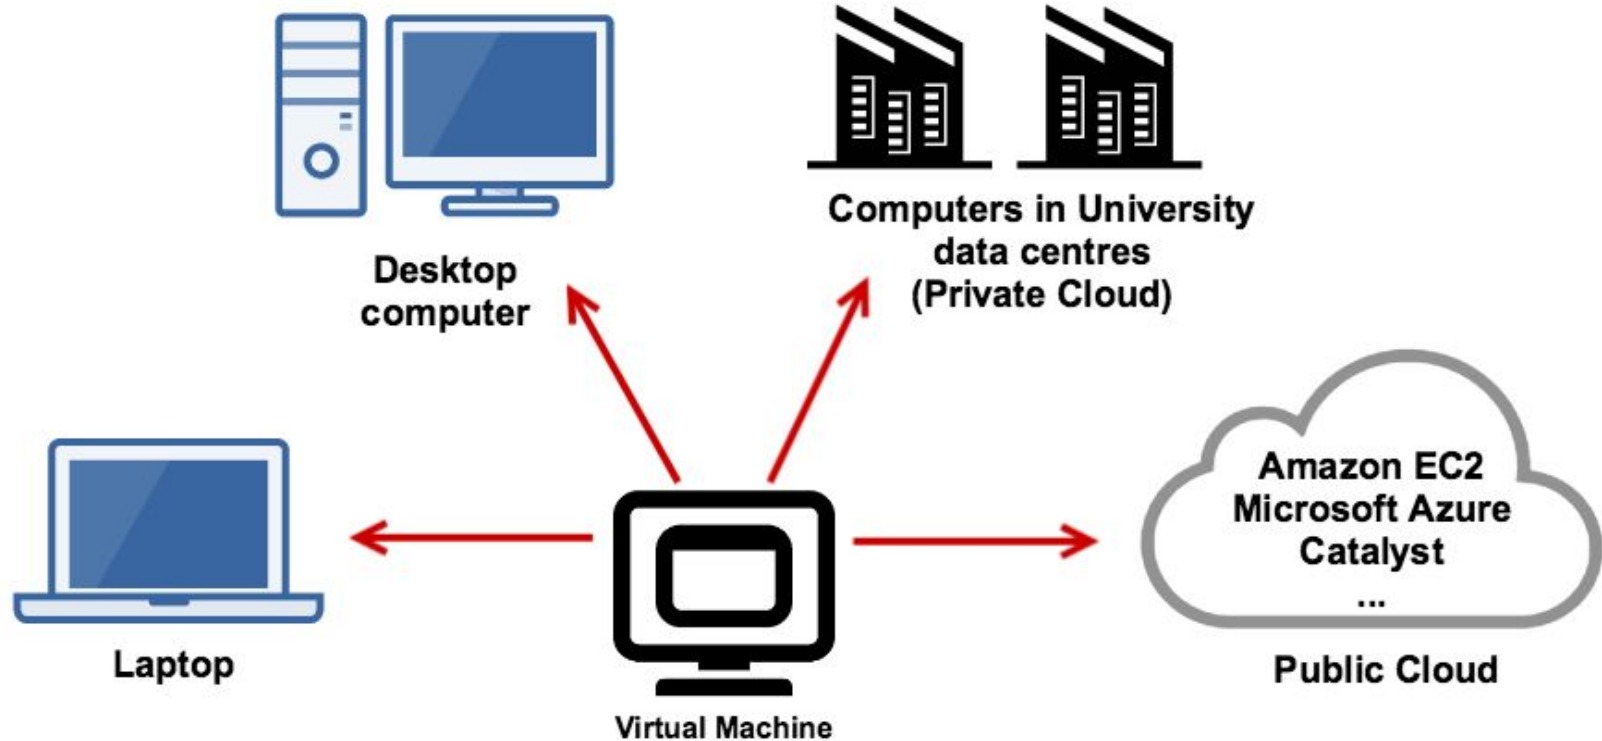
“Real machine”: physical computer

“Virtual machine” (VM): simulated computer

Differences and similarities:


- **Similarities:** Both have an operating system, network access (a real IP address), and hard disk storage.
- **Differences:** No hardware maintenance for VMs; “Snapshots” of VMs useful for backup and restore purposes; Dedicated vs shared hardware; Performance

Where do virtual machines run?



Connecting to a virtual machine

<<Demo>>



What can I do with a virtual machine?

- Run software like R, Matlab, Ansys, python, etc
- Host databases
- Host websites
- Develop software



When to use a VM over other compute resources?

PROS

- Flexibility (sizing of CPU, RAM, disk)
- Easy to request and set up
- Less maintenance
- Cost saving through shared hardware
- Collaboration

CONS

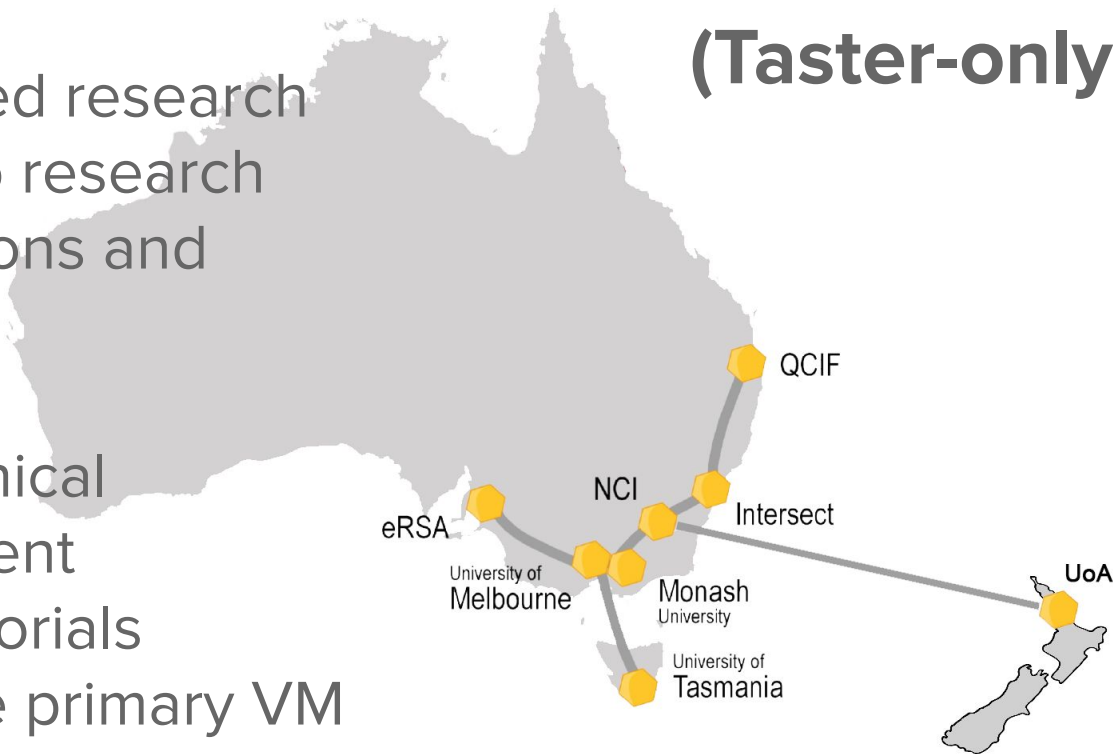
- Licensing (may be different)
- Heavy Graphics - need Graphics Processing Units (GPUs)
- Runs virtually (speed)
- Shared hardware

Virtual machine options

Nectar 1/2

- Australian federated research cloud dedicated to research
- Federated operations and support
- Self-service
- Can be quite technical
- Workshops, excellent documentation/tutorials
- Plan: Make this the primary VM service here at UoA

<<Demo>>
(Taster-only)



Nectar 2/2

- Located currently outside of the University firewall
 - More (but not infinite) freedom
 - Good for collaboration with UoA-external researchers
 - No access to data hosted on faculty drives, research drives etc. But: Plan to get Nectar on UoA internal network
- Access to GPUs
- Virtual Labs
- Advanced networking features
- Databases

Research VM Service

- Run as a collaboration between central IT and the Centre for eResearch
- Managed service: We have a consultation with you about your requirements and launch the VM for you.
- Located inside the UoA network
 - Access to faculty drives, research drives etc
 - Less easy to collaborate with UoA external researchers
- No access to GPUs
- We plan to make Nectar the only VM platform (with revised support model)

How do I request a Research VM?



Talk to us - it's free *

HackyHour (Every Thursday, 3pm, Cafe Strata)

<https://research-hub.auckland.ac.nz/#/content/1>


(*) in almost all cases

CPU OS Research-Goals
Disk-Space Backups
Software
Data-Sensitivity Who Memory
Availability

Miscellaneous about virtual machine options

- Free to researchers, within limits
- Use-cases requiring many CPU cores are directed at NeSI where possible
- Conditions of use:
 - Tell you what your research is about
 - Participation in CeR annual survey

Research Compute options - summary



Alternatives

- Public cloud: Amazon EC2, Microsoft Azure (but currently you need to spend \$\$)
- Hardware purchase

Questions?