

#### GRP

APRP - ANRP Australian National Research Platform Andrew Howard NCI Cloud Team Manager Co-Chair APAN Asia Pacific Research Platform WG



#### • NCI

#### • Bringing compute and data together

- Friction Free Data Movement
- European eXtreme Data Cloud
- Problem space
- Potential Solutions
- Research Platforms
- APRP proposed design
- Australia National Research Platform
- Data Mover Challenge
  - Australia, Singapore, Korea, Japan, USA





 The National Computational Infrastructure (NCI) is a collaboration between its foundational partners ANU, CSIRO, Bureau of Meteorology, Geoscience Australia, the major research Universities, research Institutes and Industry to provide a highly integrated computational environment for Australian researchers to enable National Research and Innovation.

nci.org.au



#### WE ENABLE AUSTRALIAN RESEARCH WITH WORLD-CLASS...

• High performance computing



#### CLOUD COMPUTING



DATA STORAGE & SERVICES













- Containers
- Role of NRENs
- Science drivers
  - Collaborative research
  - Genomics
  - Instruments like SKA











# Australian Research Clouds

- NeCTAR
  - National eResearch Collaboration Tools and Resources
  - Restructure of governance
    - Australian Research Data Commons (ARDC)
      - Combines Australian National Data Services (ANDS), Research Data Services (RDS) and NeCTAR Cloud



 The National eResearch Collaboration Tools and Resources project (Nectar) provides an online infrastructure that supports researchers to connect with colleagues in Australia and around the world, allowing them to collaborate and share ideas and research outcomes, which will ultimately contribute to our collective knowledge and make a significant impact on our society.

# **NCI** NeCTAR Cloud Operators

- NeCTAR Cloud operators
  - University of Melbourne,
  - Monash University,
  - National Computational Infrastructure (NCI),
  - Queensland Cyber Infrastructure Foundation (QCIF),
  - University of Tasmania (TPAC).



- AARNet
  - Spectrum ownership
  - Alien waves
  - Tri-versity
  - Shorter path to Europe via SingAREN and CAE-1





- What is a Research Platform
  - Notable Research Platforms
- APRP
  - Participants
    - Australia
      - NCI
      - Pawsey SuperComputer Centre
    - NRENs





- Australia National Research Platform
- How it relates to APRP
- Foundation capabilities
  - Data Movement
  - Federated authentication
  - Service orchestration
- Data Mover Challenge
  - APRP participants





- Data movement
  - File replication
  - Object replication
  - Scheduled and background transfers
- Service endpoints
  - Shared capabilities
  - Distributed data stores integrated into a single metadata namespace
- Build on advanced network capabilities



### **NCI** Friction Free Data movement

- · We need to provide our researchers with a friction free data transfer system
  - · Easy to use
  - Secure using a Federated Access system
- The network and tools should have the data in the right location at the right time
- Able to effectively use different storage tiers
  - SSD
  - Spinning Disk
  - Tape
- The researcher creates a Data Intent definition
  - Data Source
  - Data Target
    - Transfer priority (High, Medium, Low)
    - Storage performance (SSD, Disk, Tape)
  - optional Network intersection





By default TCP/IP does not perform well over high bandwidth, high delay circuits.

# A small amount of packet loss makes a huge difference in TCP performance

#### Throughput vs. Increasing Latency with .0046% Packet Loss



17

### **NC** The challenges of the Big Data era

#### How to...?

- ... orchestrate and federate Cloud, Grid and HPC [public or private] resources?
- ... Avoid software and vendor lock-in?

... overcome performance issues limiting massive adoption of virtualised Cloud resources in large data centres?

- ... exploit specialised hardware, such as GPUs or low-latency interconnections?
- ... manage dynamic and complex workflows for scientific data analysis?
- ... combine data from multiple sources and stored in multiple locations through incompatible technologies?
- ... support federated identities and provide privacy and distributed authorisation in open Cloud platforms?
- ... provide APIs to exploit the above and write applications, customisable portals and mobile views?
- ... move beyond statical location and partitioning of both storage and computing resources in data centres?
- ... distribute and deploy applications in a flexible way?
- ... exploit distributed computing and storage resources through transparent network interconnections?

### Capabilities and Requirements

- Regional connection
- Federated access
- Data capacitor capabilities
  - Local storage
- Container provisioning
  - Instantiate toolkit containers
- VM provisioning
  - Provide VM access on regionally connected DTN



- Containers
  - Docker in a well protected hosting environment
  - Singularity
    - V2
    - V3
  - Lightweight services
- Role of NRENs



- Our National Research and Education Networks are critical
  - Advanced network services
    - 100G
    - Anycast
    - IPV6
  - Data sharing services (AARNet Cloudstor)
  - National service termination point



nci.org.au







**INDIGO PaaS** 

Orchestrator



INDIGO CDMI Server













nci.org.au



- Pacific Research Platform (PRP)
  - US Initiative to build a network of Science DMZs with well tuned systems for data movement
- Asia Pacific Research Platform (APRP)
  - Regional initiative
    - KISTI Korea, NCI Australia, Perdana U Malaysia, NSCC -Singapore, Tsinghua U & NSCC Wuxi - PRC, CSIRO -Australia, Putra U - Malaysia

### NCI National Research Platform

- A collaboration between the two national HPC facilities
  - National Computational Infrastructure (Canberra)
  - Pawsey Supercomputer Centre (Perth)
- Will be connected via a dedicated 100G network path over AARNet
  - Allows for secure, high speed data transfer between the facilities
- Designed for long running services which are able to utilise HPC facilities for High Throughput



## NCI National Research Platform - projects

- KeyStone integration to support the facility specific HPC authorisation domains
- Data transfer services using Fermilab Big Data Express
  - Easy to use
  - Secure using a Federated Access system combining NCI, Pawsey and AAF namespaces
  - Supports scheduled transfers and prioritisation
  - Highly multi threaded for maximum network and storage utilisation
- Highly available national services





- Creating a national facility to support the management, processing, storage and sharing of "omics" data
- Three initial pathfinder projects
  - Zero Childhood Cancer
  - Oz Mammals
  - Genomes of Australian Plants



- NCI tender for new HPC system completed
  - A\$70M project
  - Existing HPC system will be re-purposed as additional Cloud resources for NRP and Bioscience Cloud
- Challenges
  - Disassemble and reassemble from primary to secondary data centre
  - Power and cooling
  - Networking

# **NCI** Indigo Alien Wave 300G Pawsey/NCI



# NCL APRP proposed high level design



#### National and Regional Research Platform Architecture



### Capabilities and Requirements

- Regional connection
- Federated access
- Data capacitor capabilities
  - Local storage
- Container provisioning
  - Instantiate toolkit containers
- VM provisioning
  - Provide VM access on various National Research and Commercial clouds



### High level architecture/goals

- Services may operate at a Site, National or Regional scope.
- Replication of Objects and Filesystems to support services operating in multiple Availability Zones.
- Authentication support for existing LDAP based systems and Federated identities through AAF and other federated Federations (eduGain).
- Share common best practice and personnel in design and implementation.
- Efficiently support the rapidly growing national BioInformatics activities.



- GPU access
- Object store replication
- File system replication
- Data transfer services
- Advanced Cloud development testbed
- Containers
- Message queues
- Functions

# **NCI** Conclusion

- We have started the journey
- The foundation of data movement is in progress
- Activities like the DMC are building better collaboration
- We need to investigate other shared resources



#### • For more information please contact me andrew.howard@anu.edu.au





nci.org.au



#### Acknowledgements





Australian Government Australian Research Council





Australian Government Bureau of Meteorology





Australian Government

**Geoscience** Australia