



National Computational Infrastructure

Australian National University

Data Mover Challenge reflections

Andrew Howard – Associate Director (Cloud Services & Innovation)
Chief Judge Data Mover Challenge

March 2025




Intro

Agenda

Agenda

- Introduction
- Data Mover Challenge



Challenges to large data transfer over global R&E Networks

- Disk Management (management of files for efficient transfer)
- Access Management (integration with federated services and single sign-on)
- Protocol adaptability
- OS tuning (scheduling processes. I/O strategy)
- CPU, I/O and Network latency
- Lots of small files (metadata operations)
- Parallelism
- Keeping the network full
- Dynamic adjustments based on network conditions

DMC

Data Mover Challenge

DMC Rational

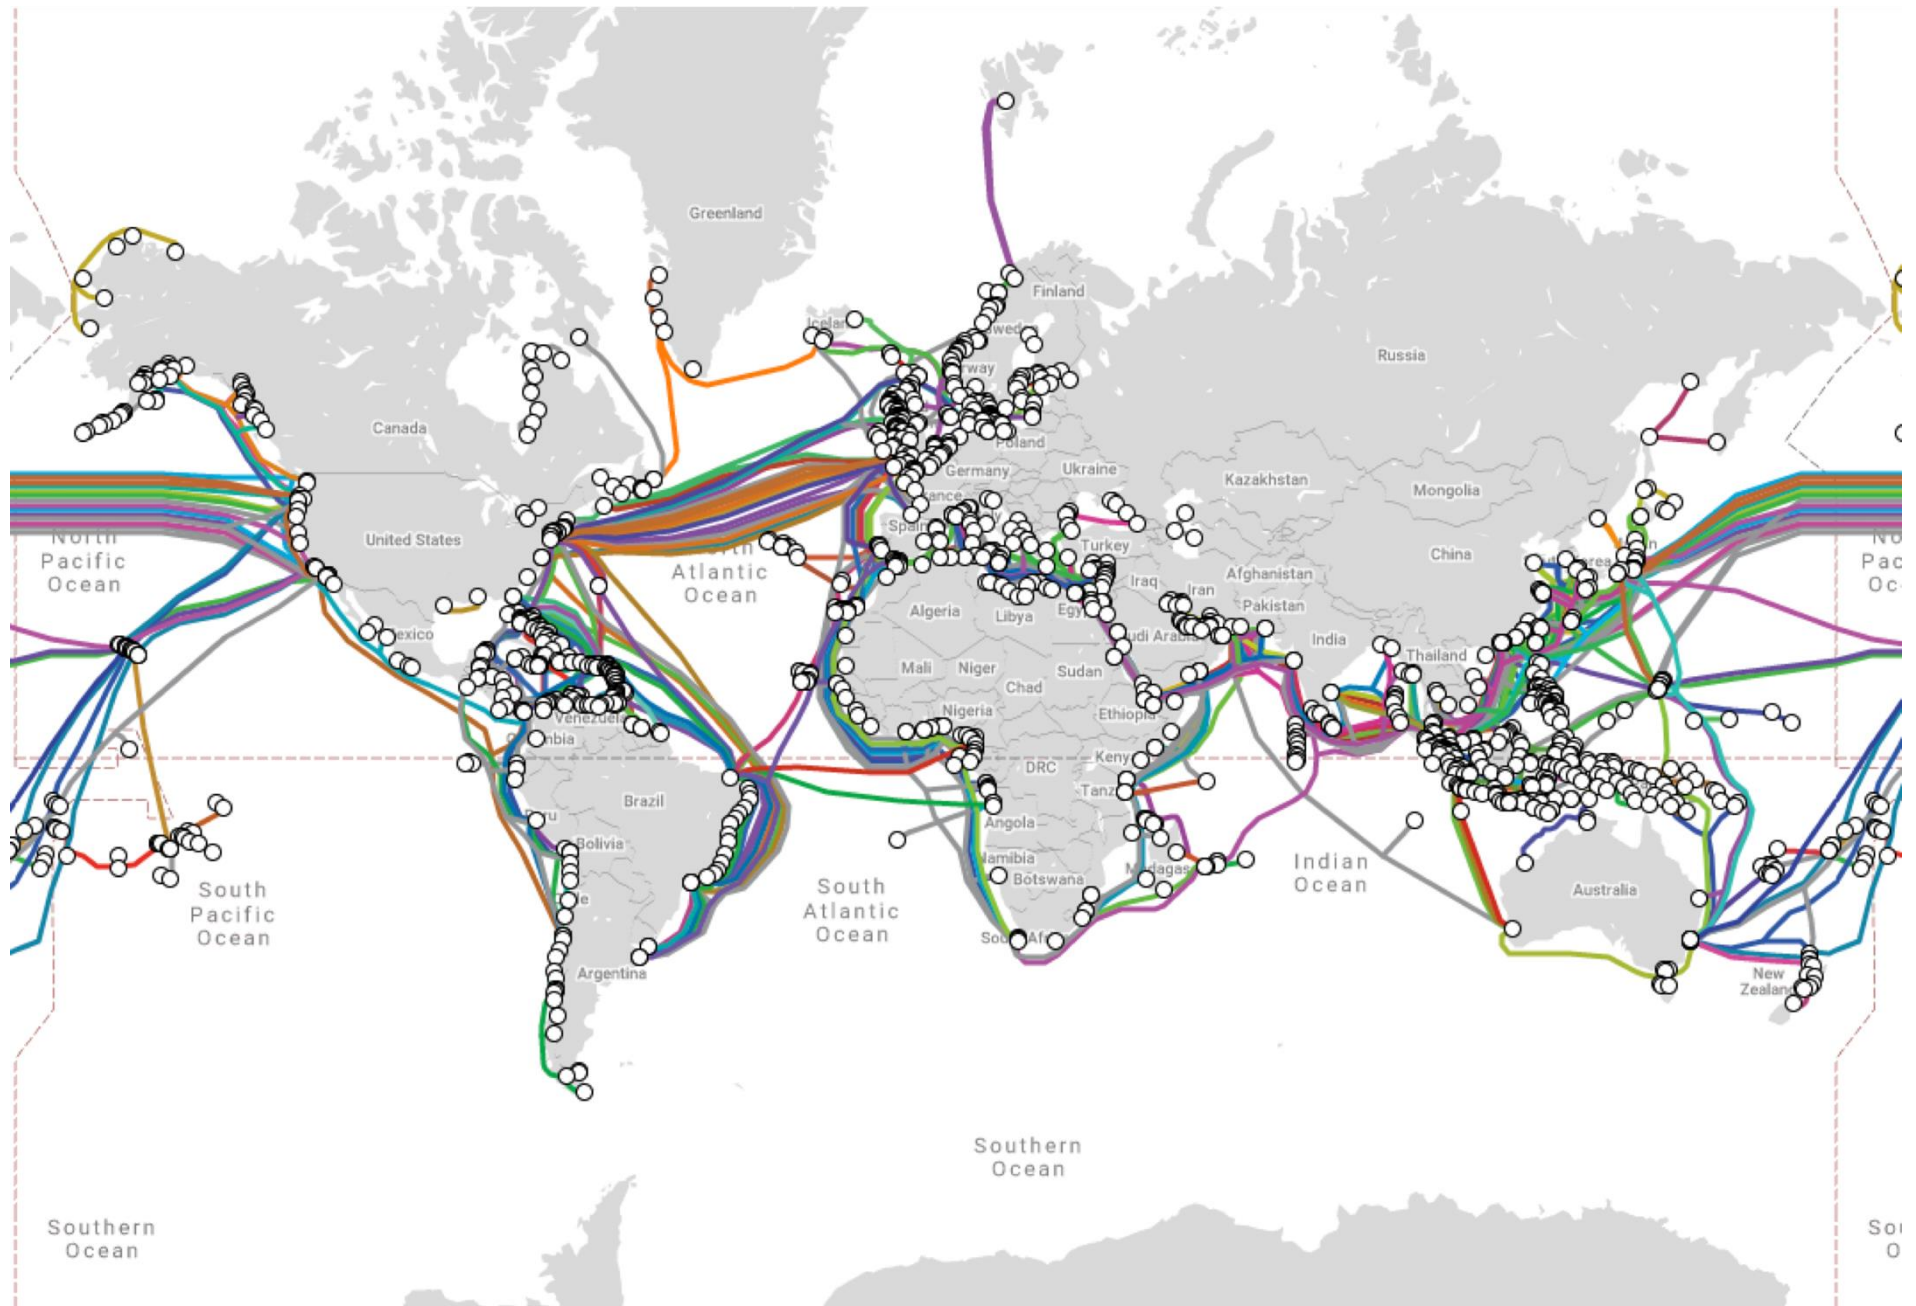
- Identify new/better methods of data transfer
 - Tools
 - Methods
- Identify emerging network tuning parameters
 - Congestion control etc
- Build relationships and trust between younger NREN engineers to reflect the trust we have between us as a community.


Data

- High performance networks are the key to data sharing
- F.A.I.R. and C.A.R.E.
- NCI does not impose ingress or egress data charges
- Friction-free data movement
 - Researchers want to access their data and don't have the time or interest to become network engineers
 - Rich range of tools

Combinative Science

- Satellites
- Telescopes
- Sensors & IoT
- Ultra high definition images
- Genome sequencers
- External data sources from around the world

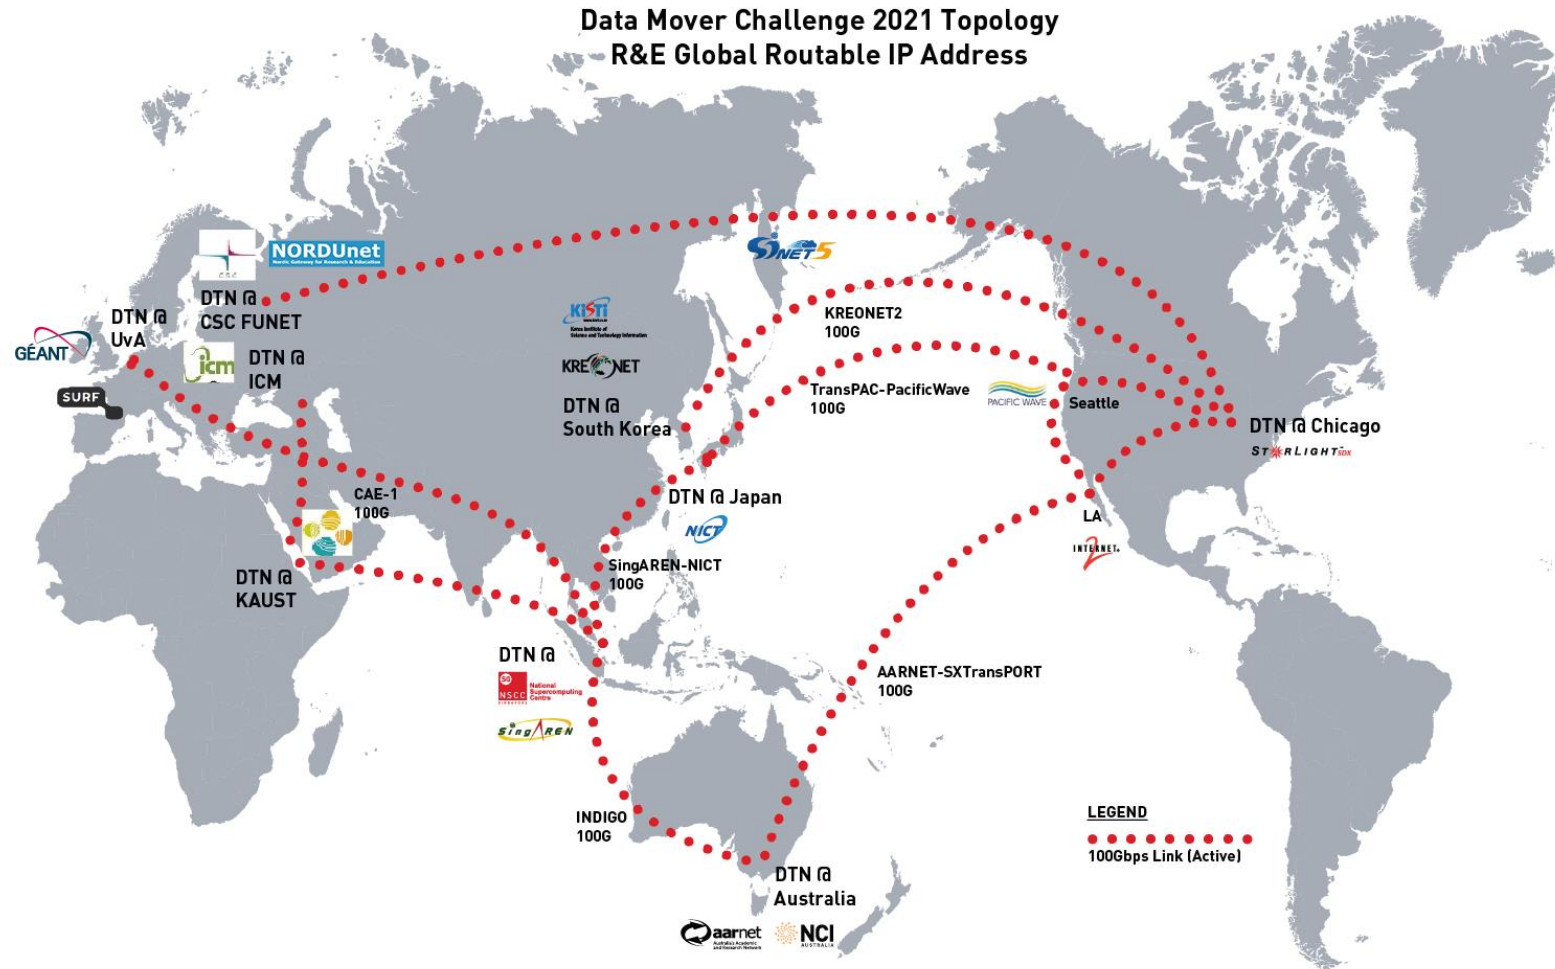




The International Data Mover Challenge

- Operates every two/three years to identify best practice for systems, networks and tools for large data transfers.
- Aims to bring together experts from industry and academia in a bid to test their software and solutions for transferring huge amounts of research data.
- Operates over production Research and Education networks.
- Critical for Australian network capabilities
 - As the most distant network user we collaborate with the International R&E community and our data suppliers like the ESA to tune for timely delivery of data products to NCI

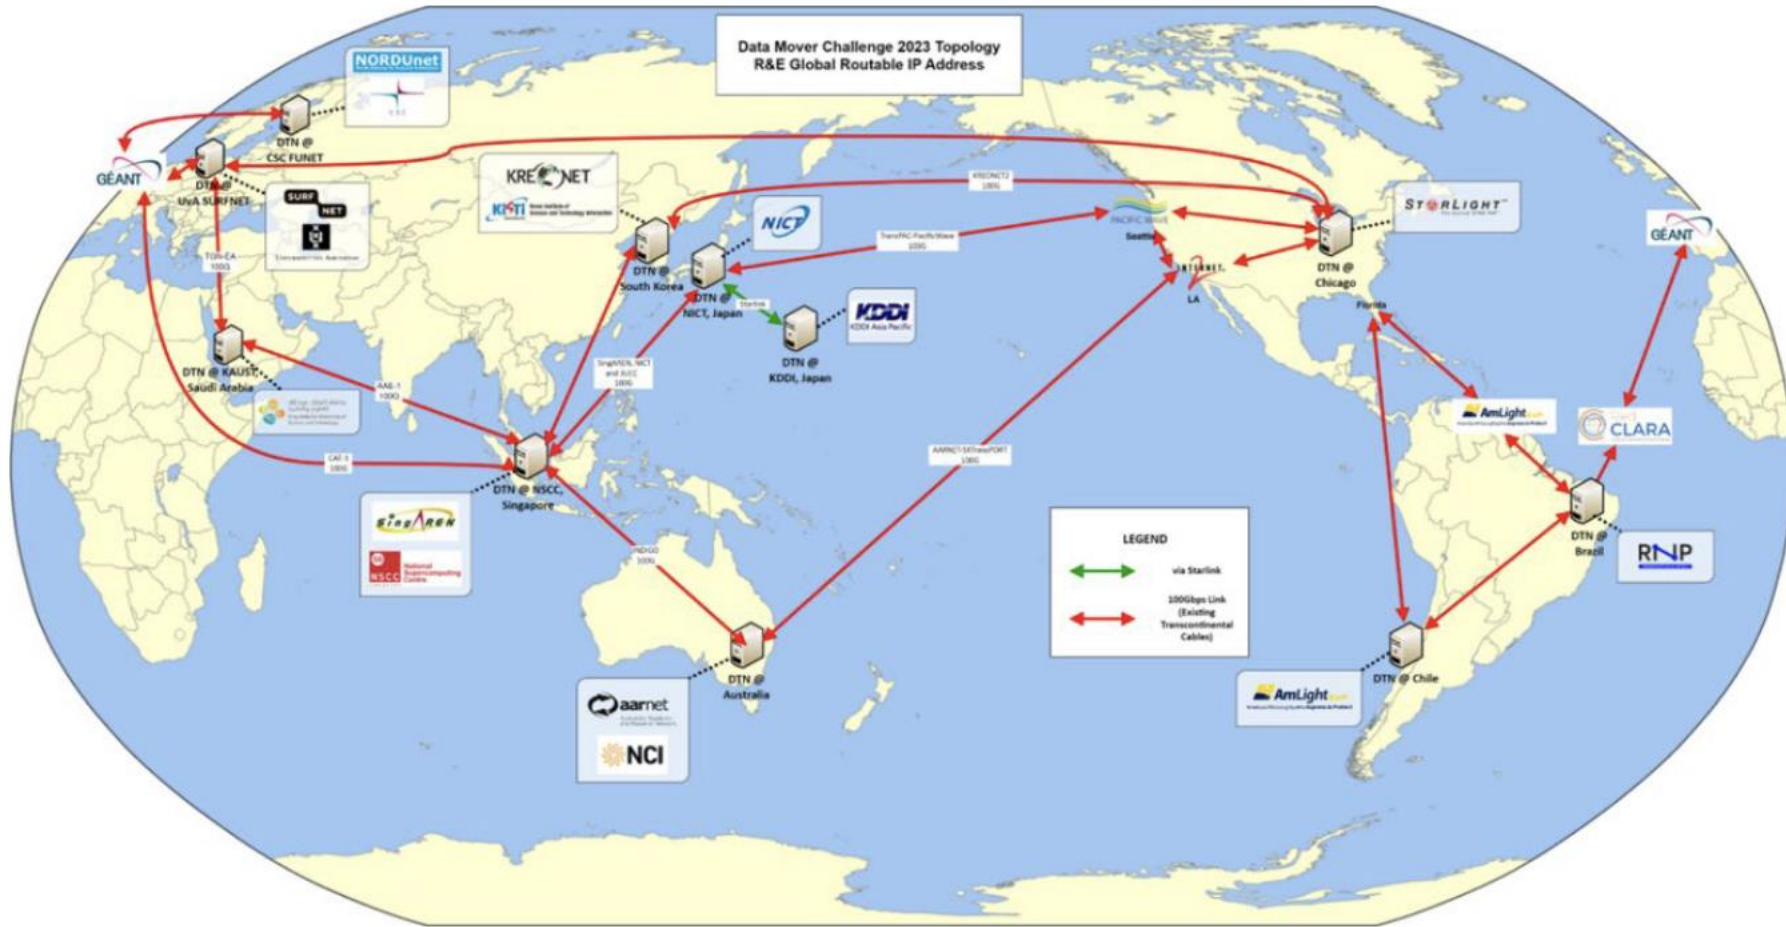
2021 DMC Network



2021

- Extended VLAN
- Richard Hughes-Jones said "this will end in tears"
- And it did
 - No traffic between DMC runs (1-2 weeks) caused ARP cache timeouts
 - Extremely difficult to troubleshoot
 - Multiple vendors
 - Multiple networks
- Work around ping regularly

2023 DMC Network



DMC23 Network

- 5 continents
 - Welcome South America
- Modelling two data transfer scenarios:
 - 100G R&E production networks
 - Starlink for field based data acquisition (thanks to KDDI)
- Australian Antarctic Division - Starlink at the South Pole (Mawson)
- Instrumentation for Judges to observe utilisation
- BIG THANKS to Simon Peter-Green and the SingAREN team for herding all the cats working with the DTN providers to clean and reset for next teams test.



DMC23 Network Partners



DMC23 DTN Partners

STARLIGHT™



ICT Solutions for Brilliant Minds



KDDI
KDDI Asia Pacific



جامعة الملك عبد الله
للعلوم والتقنية
King Abdullah University of
Science and Technology

RINIP
ORGANIZAÇÃO SOCIAL DO MCTI



UNIVERSITEIT VAN AMSTERDAM



Korea Institute of
Science and Technology Information



Challenges to large data transfer over global R&E Networks

- Disk Management (management of files for efficient transfer)
- Access Management (integration with federated services and single sign-on)
- Protocol adaptability
- OS tuning (scheduling processes. I/O strategy)
- CPU, I/O and Network latency
- Lots of small files (metadata operations)
- Parallelism
- Keeping the network full
- Dynamic adjustments based on network conditions

Judging Panel

- Prof Francis Lee - SingAREN and NTU (Singapore)
- Prof Paolo Grosso - University of Amsterdam (UVA Amsterdam)
- Hirotaka Sato - APAN-JP and KDDI (Japan)
- Ikeda-san - APAN-JP and KDDI (Japan)
- Marcos Felipe Schwarz - RNP (Brazil)
- Tim Chown - JISC (UK)
- Andrew Howard - NCI (Australia)

DMC23 Teams

- Team Falcon
- Team CiTEO
- Team Livewire
- Team Musashino
- Team Just SSH
- Team Raysync x Robust HPC Alliance

Team Falcon

- University of Reno Nevada - USA
- PhD project
- Inspired by hands on experience
 - Understanding the data provides insight into the best methods for data transfer
 - Lots of small files (metadata operations)
 - Large files (parallel transfer)
 - Pre-profile data set
 - Heuristic feedback between sender and receiver agents
 - Dynamic
 - I/O optimisation
 - Network parameter adjustment

Team CiTEO

- DTN Optimization using Dynamic Network Estimation and DTN Impedance matching.
 - CIENA environment for network innovation, network and testbeds.
- Automation and AI to optimise transfers.
 - Homogeneity of DTN deployments
 - Switch buffer size, hosts buffer sizes
 - Timely delivery (adaptive during transfer)
 - Adapts to non deterministic network conditions
- Provided as DTN as a Service (DTN-AAS) layered on Jupyter notebooks

Team Livewire

- Arcitecta - Commercial - Australia
- Mediaflux product - Livewire service (Big Data Through the Eye of a Needle)
 - Mediaflux provides a complete fabric for data management and transfer
 - Metadata driven automated replication, duplication and distribution
 - Parallel streams (default 10)
 - Visualisation of data flows for rapid performance indication
 - Virtualised storage layer for optimised I/O handling

Team Musashino

- University of Tokyo - Japan
- High-performance and Flexible Protocol and its application challenges high-speed file transfer in 100G on global LFNs
- HPfp - Protocol designed to perform well in impaired network conditions
 - Natural Disasters in Japan
 - Variety of possible network services which can change dynamically
 - UDP based designed for Long Fat Network with packet losses
 - Socket library based implementation
 - Modes (fair, modest, aggressive and more)
 - Performed extremely well over Starlink (during a week of torrential rains and storms)

Team Just SSH

- Ryo Nakamura - PhD project
- <https://github.com/mscp>
- Client side modifications of scp
 - Multiple streams of sftp
 - Immediate 10x performance improvement over standard scp
 - Able to adjust and tune parameters based on observed performance (defaults are adequate but additional tuning depends on network conditions)

Team Raysync x Robust HPC Alliance

- Raysync - Commercial - Malaysia, China
- UDP based adaptive tuning protocol
 - *Appears to be a reverse engineering of Aspera*
- Excels in “degraded” network conditions
 - Adapts to Firewalls
 - Compression
 - Effective for high throughput transfers from the field
- Complete end to end data management platform

Thanks to everyone

- Collectively we are doing amazing things.
- We need to bring the next generation of engineers along
 - SCInet is an excellent example
 - Building the trust we collectively have in the next generation
- It's a pleasure to work together with you all.
- See you all next conference !



Questions ?



NCI
AUSTRALIA

NCI Contacts



General enquiries: +61 2 6125 9800




Support: help@nci.org.au



Email: andrew.howard@anu.edu.au

Email: daniel.rodwell@anu.edu.au



Address

NCI, ANU Building 143

143 Ward Road

The Australian National University

Canberra ACT 2601