


# 'Move That Data!' Data Mover Challenge Judging Reflections

Andrew Howard

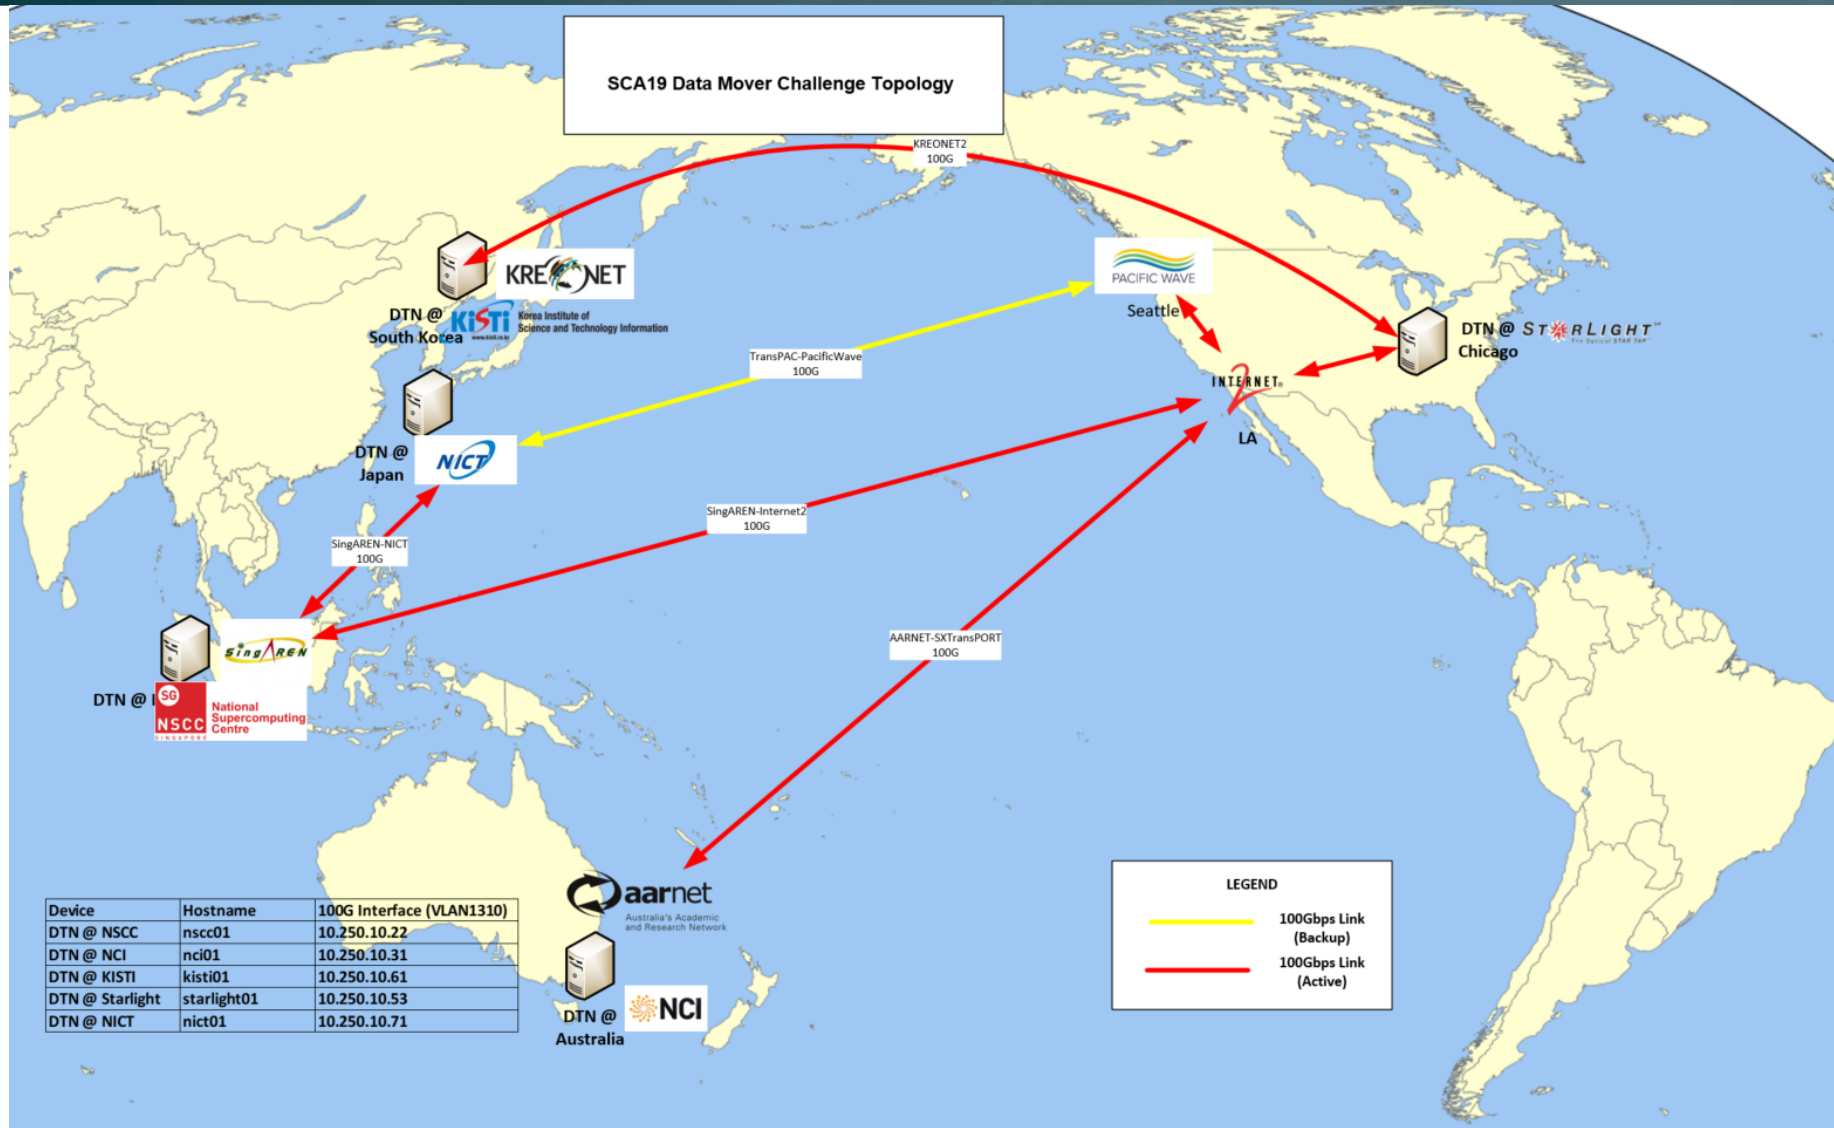


**NCI**  
AUSTRALIA



**National  
Supercomputing  
Centre**

- DMC
- Challengers
  - SEAIP/NCHC+Starlight
  - Fermilab
  - Zettar Inc
  - The University of Tokyo
  - Argonne National Laboratory
  - iCAIR/Northwestern University
  - JAXA/Fujitsu



- Succession planning
  - Next generation of network engineers
  - NOCs
- Build trust networks
- Containerised transfer tools
  - Pull the desired toolkit from repo to each DTN
  - Version control
- Network capability validation



- Most participants requested
  - NIC ring buffers (Mellanox)
  - Pacing
  - NUMA socket affinity

- Bring the community together
- Testing between 100G connected HPC centres (NSCC, NCI, NICT, KISTI, StarLight)
- Build a portfolio of containersied toolkits tested within the platform
- Create the best practice for system and network tuning in the Asia Regional context



- Production networks
- Existing DTNs
- Local DTN storage only (disk I/O bottleneck)
- No sudo (root) access
- Workflows and tools must be installed in Singularity containers
- Data set 2Tb with 1431 files built from actual data (Genomics, Climate, Video) “Lots of small files”
- VLAN does not allow monitoring of the outer production network congestion

- Brings together experts from industry and academia, in a bid to test their software across servers located in various countries (Australia, Japan, Singapore, USA) that are connected by 100G international networks.
  - Each team given 1 week to deploy the software on globally distributed nodes.
  - The transport protocol must minimally support a data transfer rate of 10Gbps.



- Memory to Memory
- Disk to Disk
  - Maximum sustainable transfer rate from disk to disk
  - Usability of the software
  - Cost
  - Scalability
- Innovation

- ESNet
  - A Scalable Network Design Pattern for Optimising Science Data Transfers
    - The Science DMZ
      - is a portion of the network, built at or near the campus or laboratory's local network perimeter that is designed such that the equipment, configuration, and security policies are optimised for high-performance scientific applications rather than for general-purpose business systems or “enterprise” computing.
      - model addresses common network performance problems encountered at research institutions by creating an environment that is tailored to the needs of high performance science applications, including high-volume bulk data transfer, remote experiment control, and data visualisation.

- A network architecture explicitly designed for high-performance applications, where the science network is distinct from the general-purpose network.
- The use of dedicated systems for data transfer.
  - Data Transfer Node (DTN)
- Performance measurement and network testing systems that are regularly used to characterise the network and are available for troubleshooting.
- Security policies and enforcement mechanisms that are tailored for high performance science environments





- Each team will leverage on the global Research and Education infrastructure for the Data Mover Challenge.
- The deployed DTNs will connect all locations, including Australia, Japan, Singapore and USA.
- Teams must determine the connection type that will be used, for example, point-to-point, multipoint or distributed.
- Multiple Virtual Local Area Networks (VLANs) will be set up, per location, to provide inter-connectivity.
- Network monitoring will be set up to capture data transfer statistics.



**National  
Supercomputing  
Centre**

- AU
  - Andrew Howard (NCI)
- SG
  - Alan Davis (NSCC)
  - Prof. Lawrence Wong (NUS)
  - Prof. Francis Lee (NTU)
- US
  - Eric Pouyoul (ESNet)

- Thanks to:
  - Alvin Chiam (NSCC)
  - Simon Peter Green (SingAREN)
  - John Foo (SingAREN)



- Participants

- SEAIP/NCHC+Starlight (Week of: 14-Jan-19)
- Fermilab (Week of: 28-Jan-19)
- Zettar Inc (Week of: 4-Feb-19)
- The University of Tokyo (Week of: 11-Feb-19)
- Argonne National Laboratory (Week of: 18-Feb-19)
- iCAIR/Northwestern University (Week of: 25-Feb-19)
- JAXA/Fujitsu (Week of: 4-Mar-19)



- SEAIP DTN-as-a-Service
  - Jupyter Notebook controller and user interface
  - NUTTCP transfer tool
- Team Members
  - Team Lead: Steven Shiau, Co-Leads: Jim Chen, Te-Lung Liu
  - Participants: Weicheng Huang, Ceasar Sun, Thomas Tsai, WeiYu Chen, Hui-Lan Grace Lee, Jen-Wei Hu, Li-Chi Ku, Se-Young Yu, Fei Yeh, Xiao Wang, Dinh Van Dzung, Sheldon Knuth, Denny Hermawan, Heru Suhartanto, Kusmardi Kusmardi, Nur Maya, Manaschai Kunaseth, Susumu Date, Sri Wahjuni, Chalernpol Charnsripinyo.

- BigData Express



- Team Members

- Fermilab: Wenji Wu, Qiming Lu, Liang Zhang, Sajith Sasidharan, Phil DeMar
- StarLight: Jim Chen, Joe Mambretti, Se-young Yu, Fei Yeh
- KISTI: Jin Kim, Seo-Young Noh
- KREONET: Buseung Cho, Chanjin Park

- Zettar zx hyperscale data distribution software platform
- Team Members
  - Chin Fang, Ph.D. Founder & CEO
  - Igor Solovyov, MSCS, Founding Principal Engineer
  - Alex Nazarenko, MSCS, Founding Senior Engineer





- Secure Data Reservoir
- Team Members
  - Junichiro Shitami, Goki Honjo, Kei Hiraki, Mary Inaba



- Using GridFTP and Globus Online for Large Data Transfers
- Team Members
  - Joaquin Chung, Zhengchun Liu, Tekin Bice, Rajkumar Kettimuthu, and Ian Foster



- STARLIGHT DTN-as-a-Service for Intensive Science
  - NUMA, NVMe and RAID configuration module
  - GNU/Linux and 100GbE NIC optimisation module
  - A high-speed parallel data transfer module
  - OpenNSA network provisioning module
  - Monitoring and graphing module
- Team Members
  - Se-young Yu, Fei Yeh, Xiao Wang, Jim Chen

- Smart Communication Optimizer
  - Proprietary protocol to accelerate standard tools
- Team Members
  - JAXA: Naoyuki Fujita, Hirofumi Ohkawa, Hidekazu Mikai
  - Fujitsu: Yoshio Sakaguchi, Sho Kato, Kazuhiro Miyashita, Hiroshi Takamura, Yuta Kawamura





- For more information please contact me  
**[andrew.howard@anu.edu.au](mailto:andrew.howard@anu.edu.au)**

# Questions ?



