



A*STAR/NSCC/SingAREN

Key national and International initiatives

Yves Poppe
A*STAR Computational Resource Centre, Singapore

September 28th 2016, Miami, FL
16th annual GLIF Workshop

Singapore



Physical Land Area: 718 sq km

Population: 5.5 million

3.9 million (70.8%)

Singapore Residents

Literacy Rate: 96.7% (Aged 15 & above)

76.4% of those aged 25-34 years
have tertiary qualifications

2014 GDP S\$390.1bn (US\$307.9bn)

Real Growth : 2.9%

Per Capita GDP: S\$71,318
(US\$56,284)

Independent since 9 August 1965

Sources: *Singapore in Brief 2015, Department of Statistics*

Agency for Science, Technology and Research

>1,700

INDUSTRY PROJECTS A YEAR

5 INDUSTRY PROJECTS A DAY

>2,800

PAPERS PUBLISHED A YEAR

>200

LICENSES A YEAR

4 LICENSES A WEEK

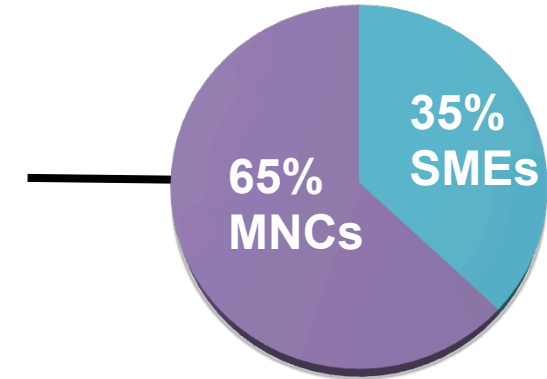
1

RSE SPUN OUT TO INDUSTRY A DAY*

** average number of Research Scientists and Engineers (RSE) per working day in a calendar year*

>240

PATENTS FILED A YEAR



>13

START-UPS A YEAR

A*STAR Computational Resource Centre

Scope:

- serve A*STAR users
- exploration of novel technologies and architectures

Many unique systems:

- Cumulus (IBM), Cirrus (IBM), Aurora (SGI), Fuji (Fujitsu), Axle (HP)

Large range of architectures:

- Intel x86, IBM Power, NVIDIA Tesla, AMD FirePro, Intel Xeon Phi, Micron Automata

Involvement in exploration projects:

- InfiniCortex – InfiniBand ring around the world
- Intel Center of Excellence
- IBM OpenPower Foundation member (*4 x IBM Power 8 servers*)
- Micron Automata Center of Excellence (*2 x Automata boards*)

National Supercomputing Centre (NSCC)



~1 PFLOP System

- **1,288 nodes** (dual socket, 12 cores/CPU E5-2690v3)
- **128 GB DDR4 RAM/ node**
- **10 Large memory nodes** (1x6TB, 4x2TB, 5x1TB)

FUJITSU



13PB Storage

- **HSM Tiered, 3 Tiers**
- **I/O 500 GB/s flash burst buffer**
- **10x Infinite Memory Engines (IME)**



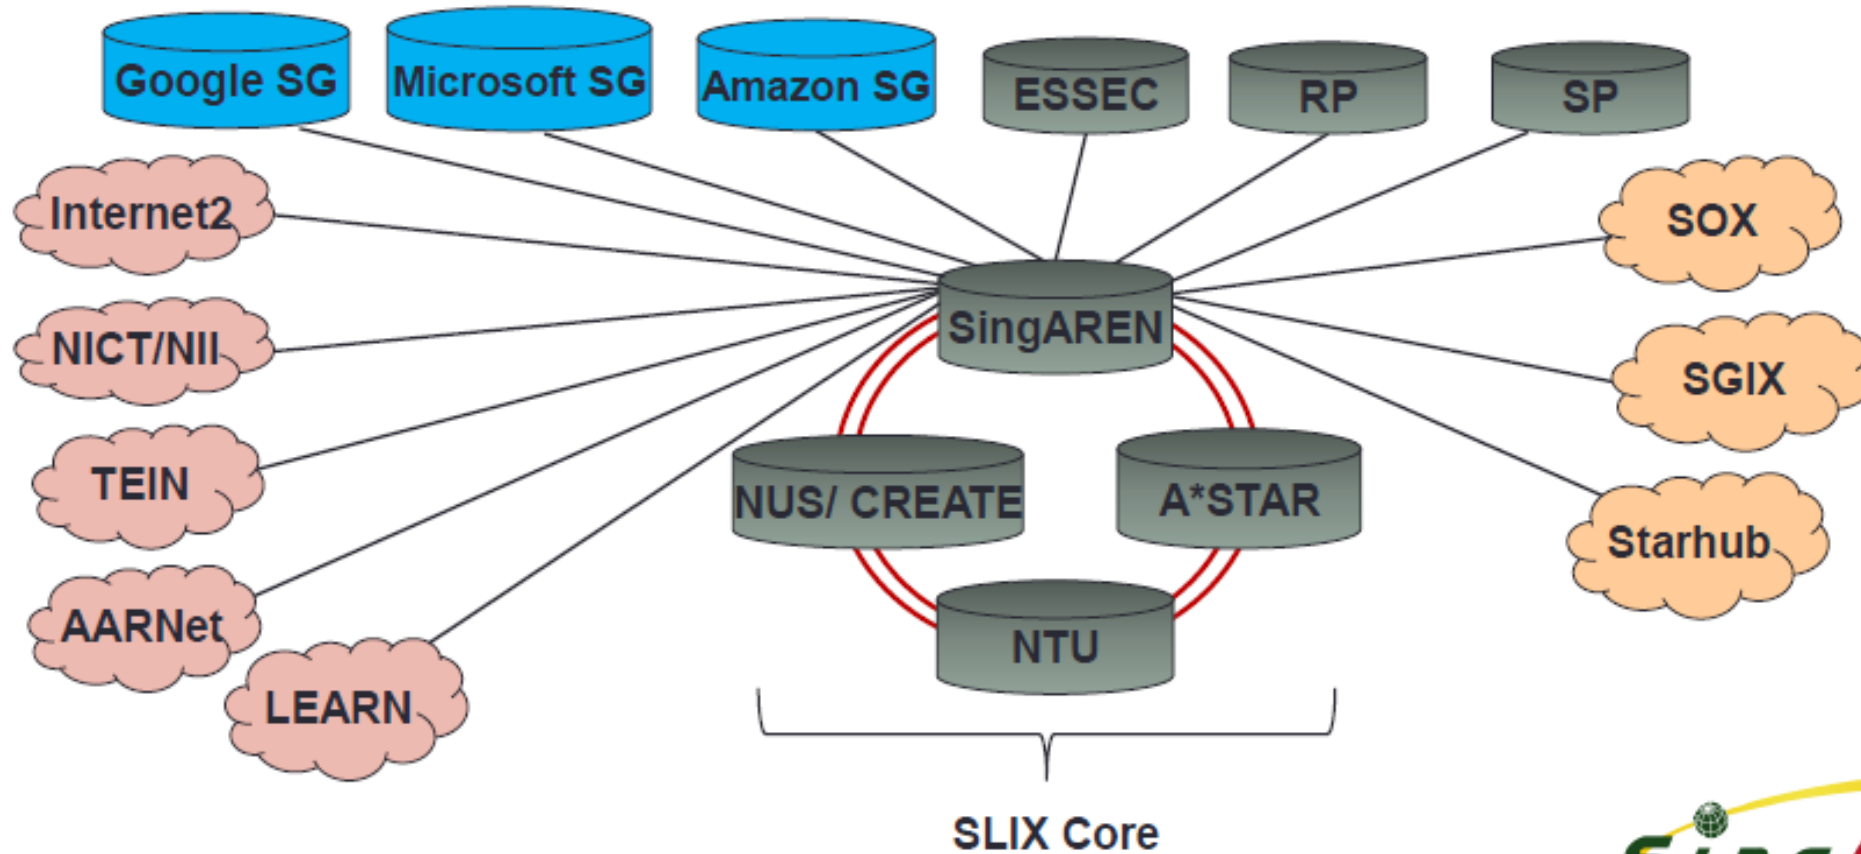
EDR Interconnect

- **EDR (100Gbps) Fat Tree within cluster**
- **InfiniBand connection to remote login nodes at stakeholder campuses (NUS/NTU/GIS) at 40/80/500 Gbps throughput**

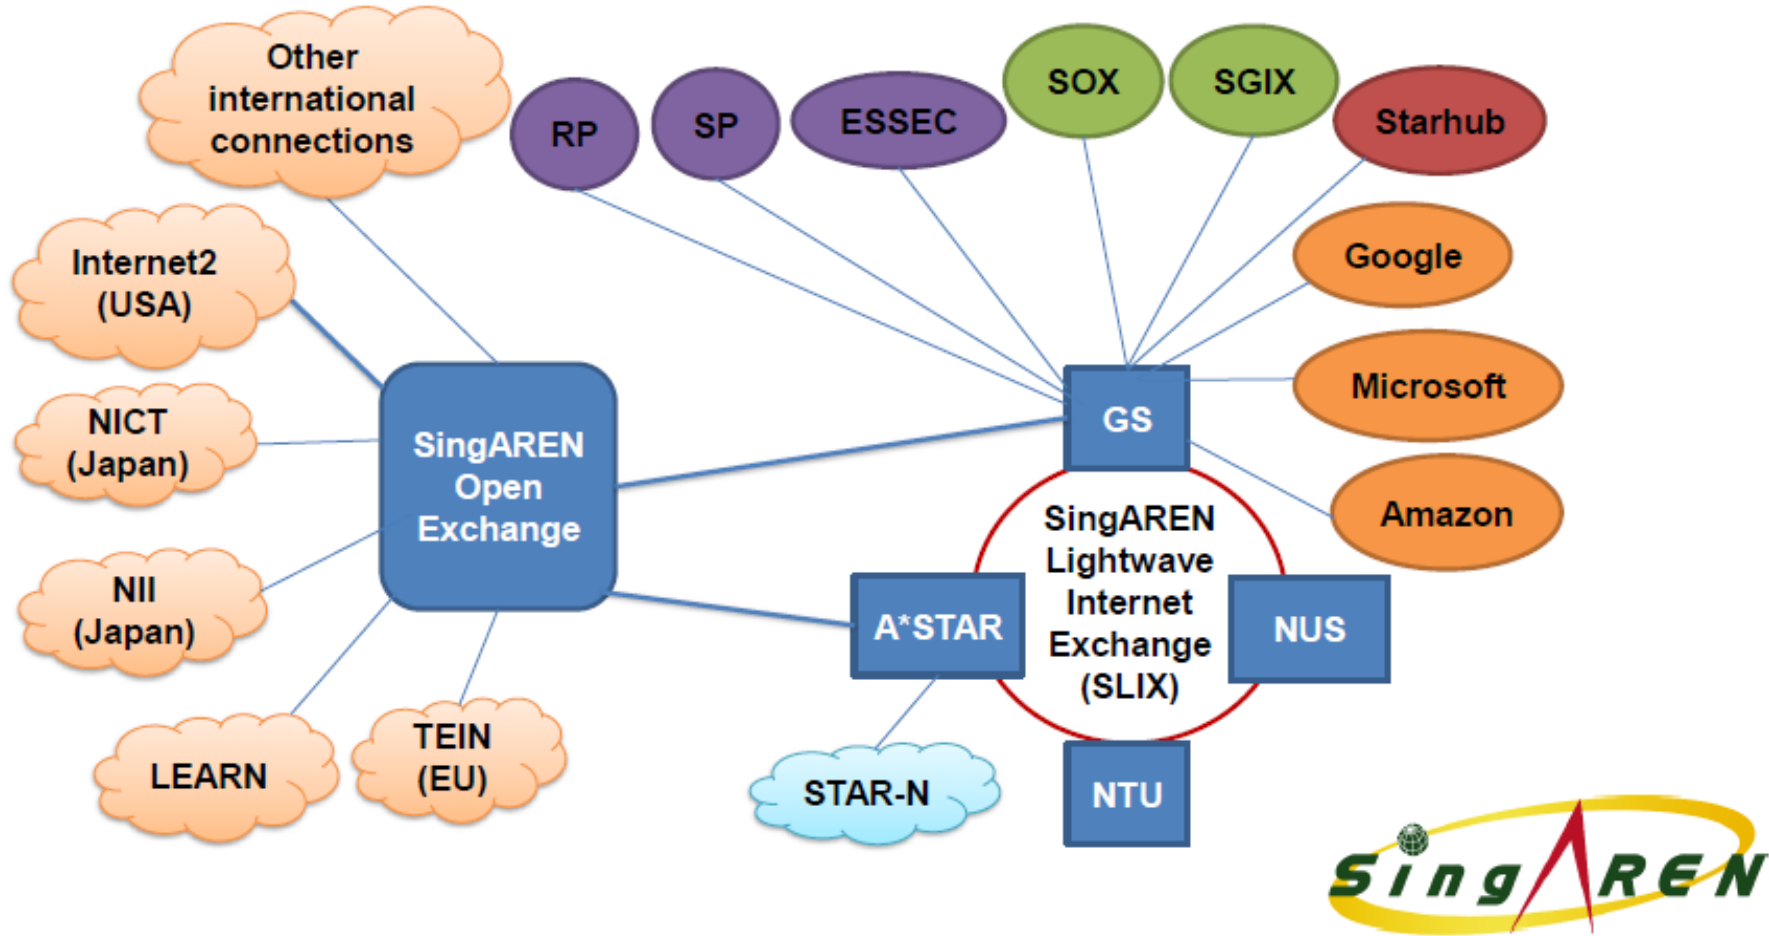
Mellanox
TECHNOLOGIES



SLIX: SingaREN Lightwave Internet Exchange



SingaREN Open Exchange



'Maximize my effective throughput between my storage and my compute'

Fusionopolis and Biopolis are 2km apart



Tests started with Mellanox METRO-X early 2013. Today the sites are connected with nx40gbps connections running native InfiniBand and reaching approx. 98.4% of maximum theoretical possible throughput. Dark fibre running at 400Gbps, soon 1Tbps using Infinera connects the two sites.

RDMA - InfiniBand – RoCE: Going the distance

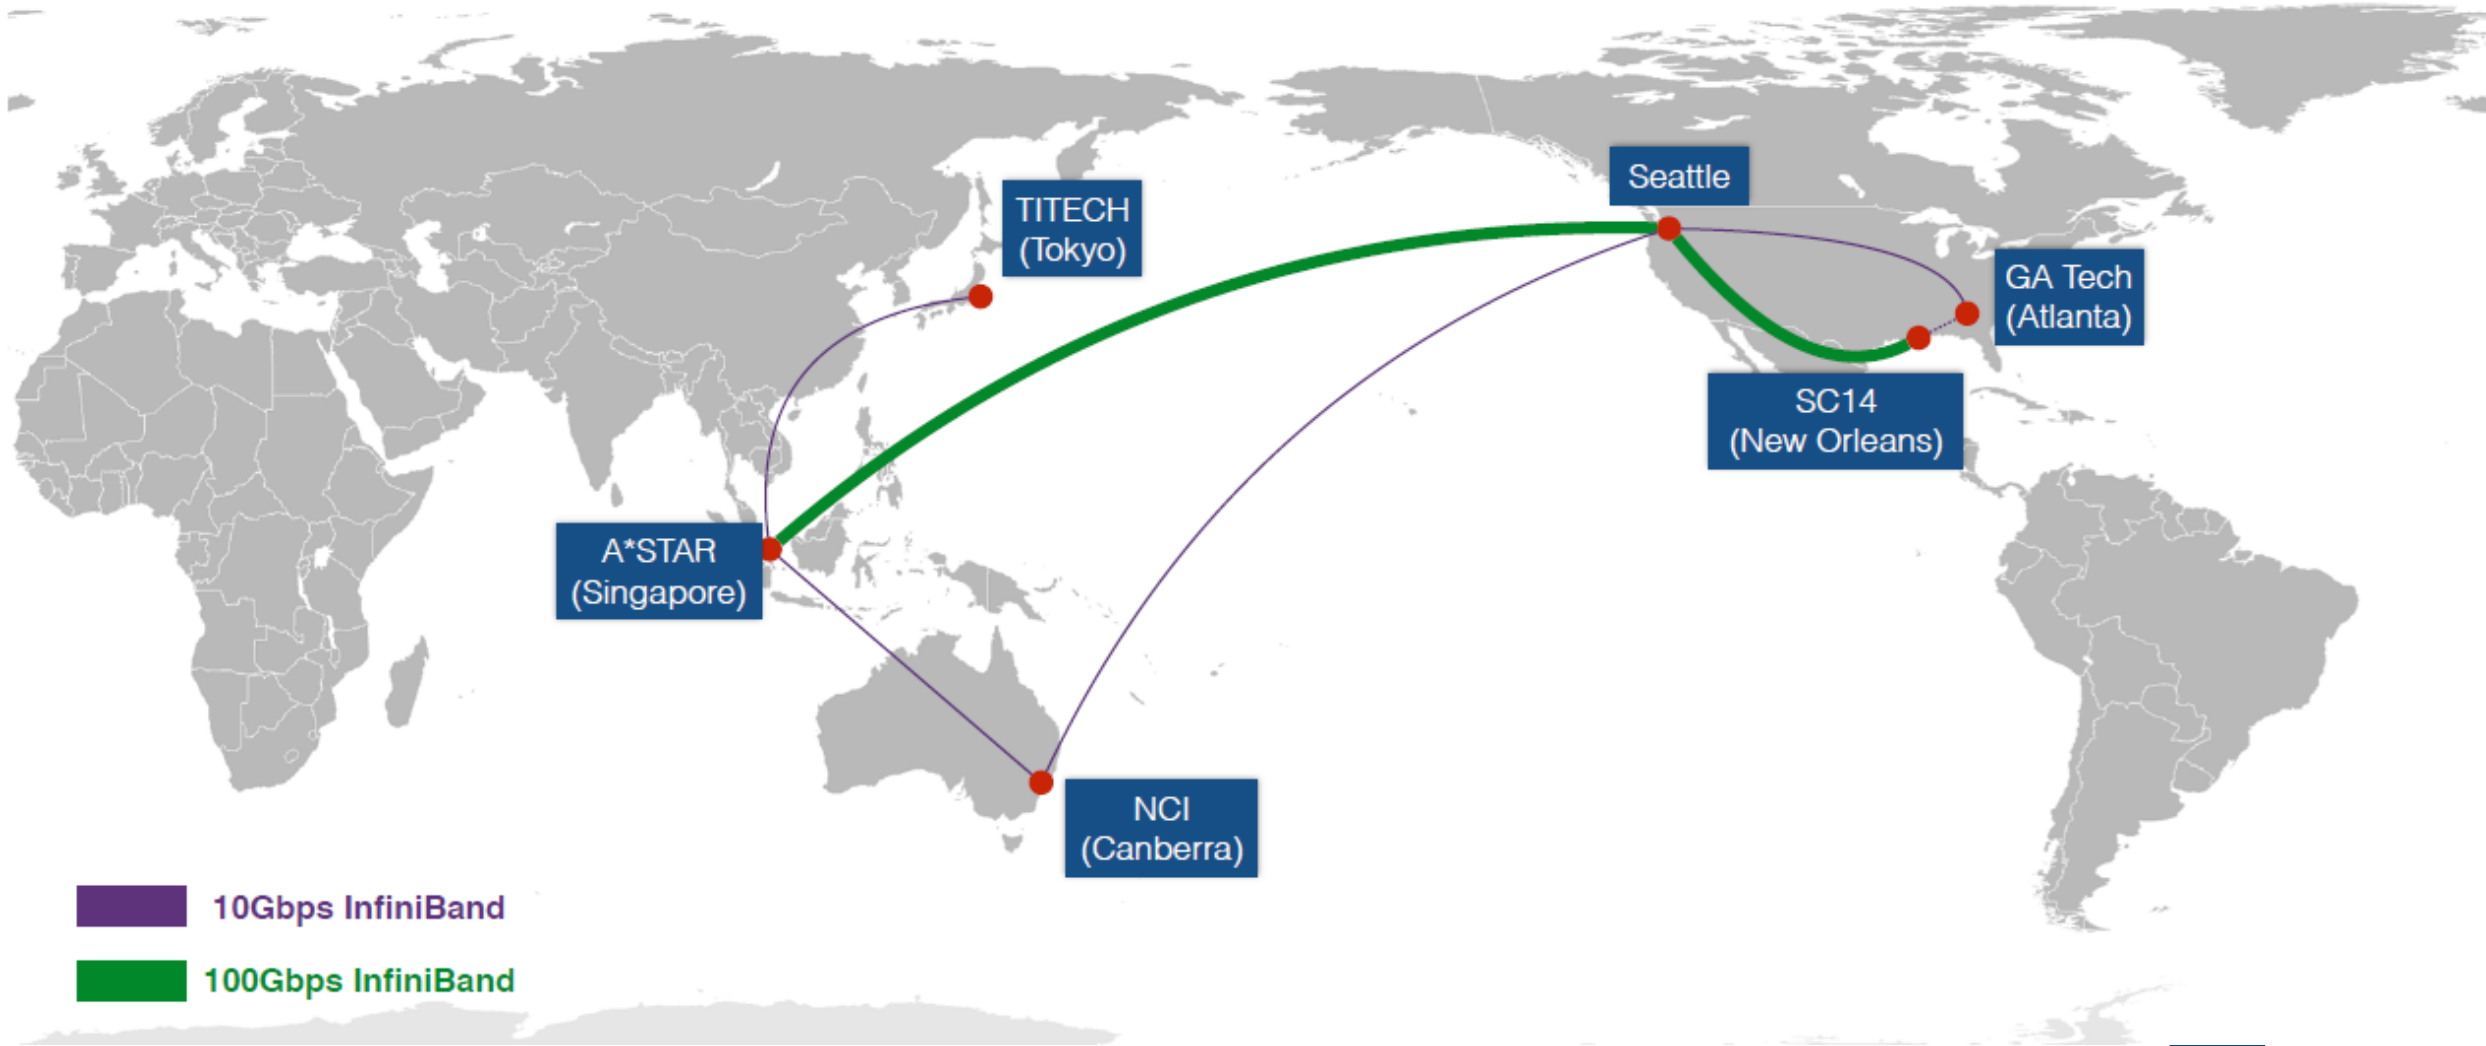
- Remote Direct Memory Access (RDMA) is the core concept for efficient operation between the multiple cores of a cluster.
- Ethernet, InfiniBand and proprietary standards are used for communication between cores. InfiniBand has become the protocol of choice in recent years providing the best throughput at the lowest latency and runs on more than 50% of the Top100 supercomputers. The Ethernet camp has developed RoCE (RDMA over converged Ethernet) in response.
- A long distance version of InfiniBand was developed under the impetus of the US military resulting in two current suppliers, Bay Networks and Obsidian Strategics. Considering evaluation of long distance implementation of RoCE.



Singapore's approach to exascale: InfiniCortex

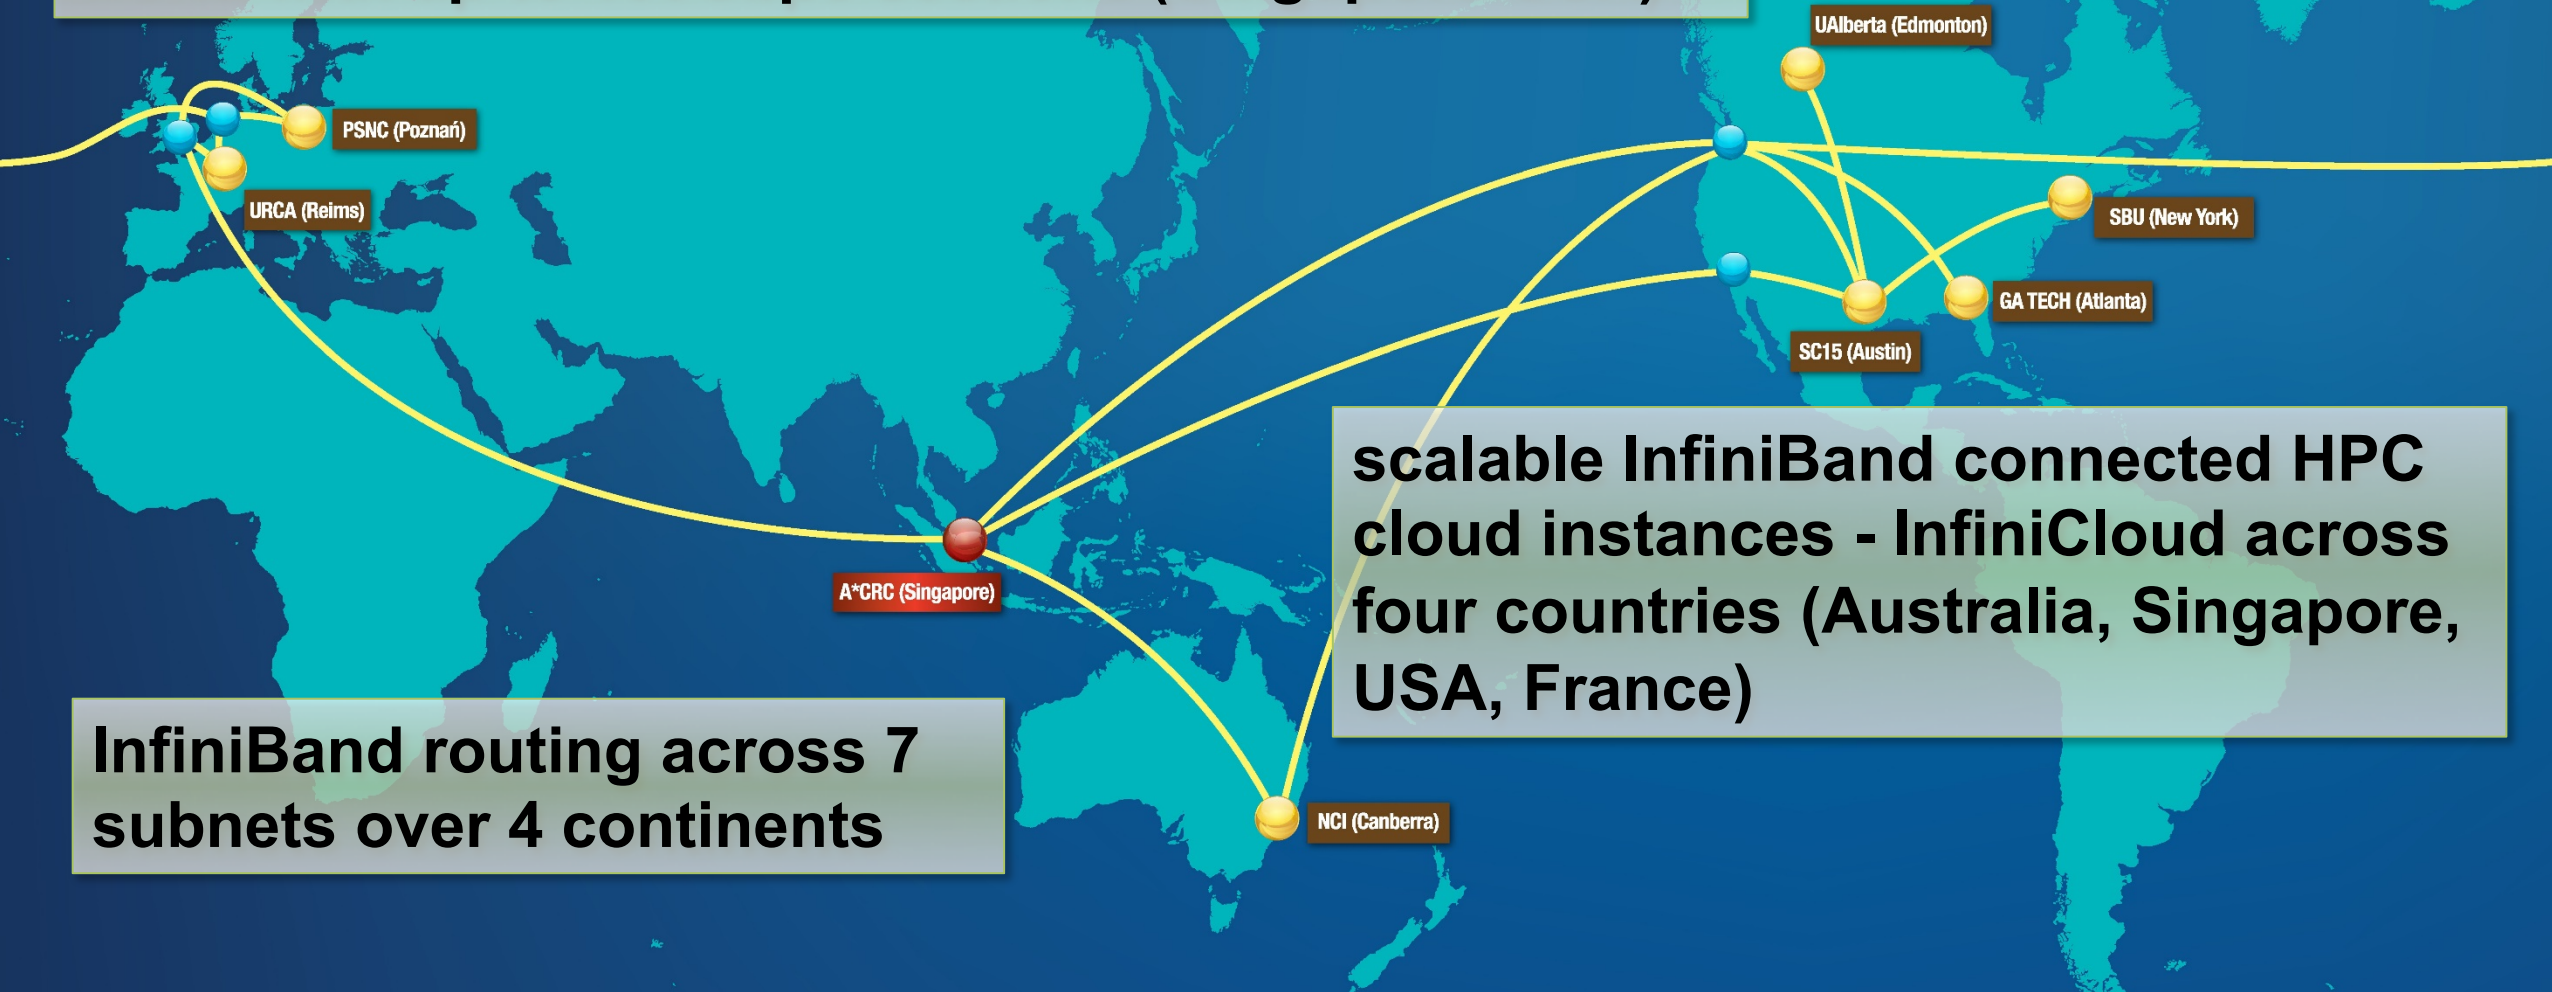
- A geographically dispersed constellation of compute, storage and associated power needs, working as one; not grid, not cloud.
- The five elements needed to succeed are lined up:
 - Supercomputer interconnect topology based on graph theory work done by Y Deng, M. Michalewicz and L. Orłowski.
 - Availability of high speed uncongested bandwidth
 - Long distance InfiniBand to increase effective throughput over any given link and InfiniBand routing.
 - Application layer: from simple file transfers to complex workflows with Oakridge developed ADIOS and multi-scale models.
 - Partnerships: Critical mass is being reached.

InfiniCortex at SC 2014



InfiniCortex at SC 2015

largest spanning InfiniBand network: ring-around-the-world with up to 100Gbps sections (Singapore-USA)



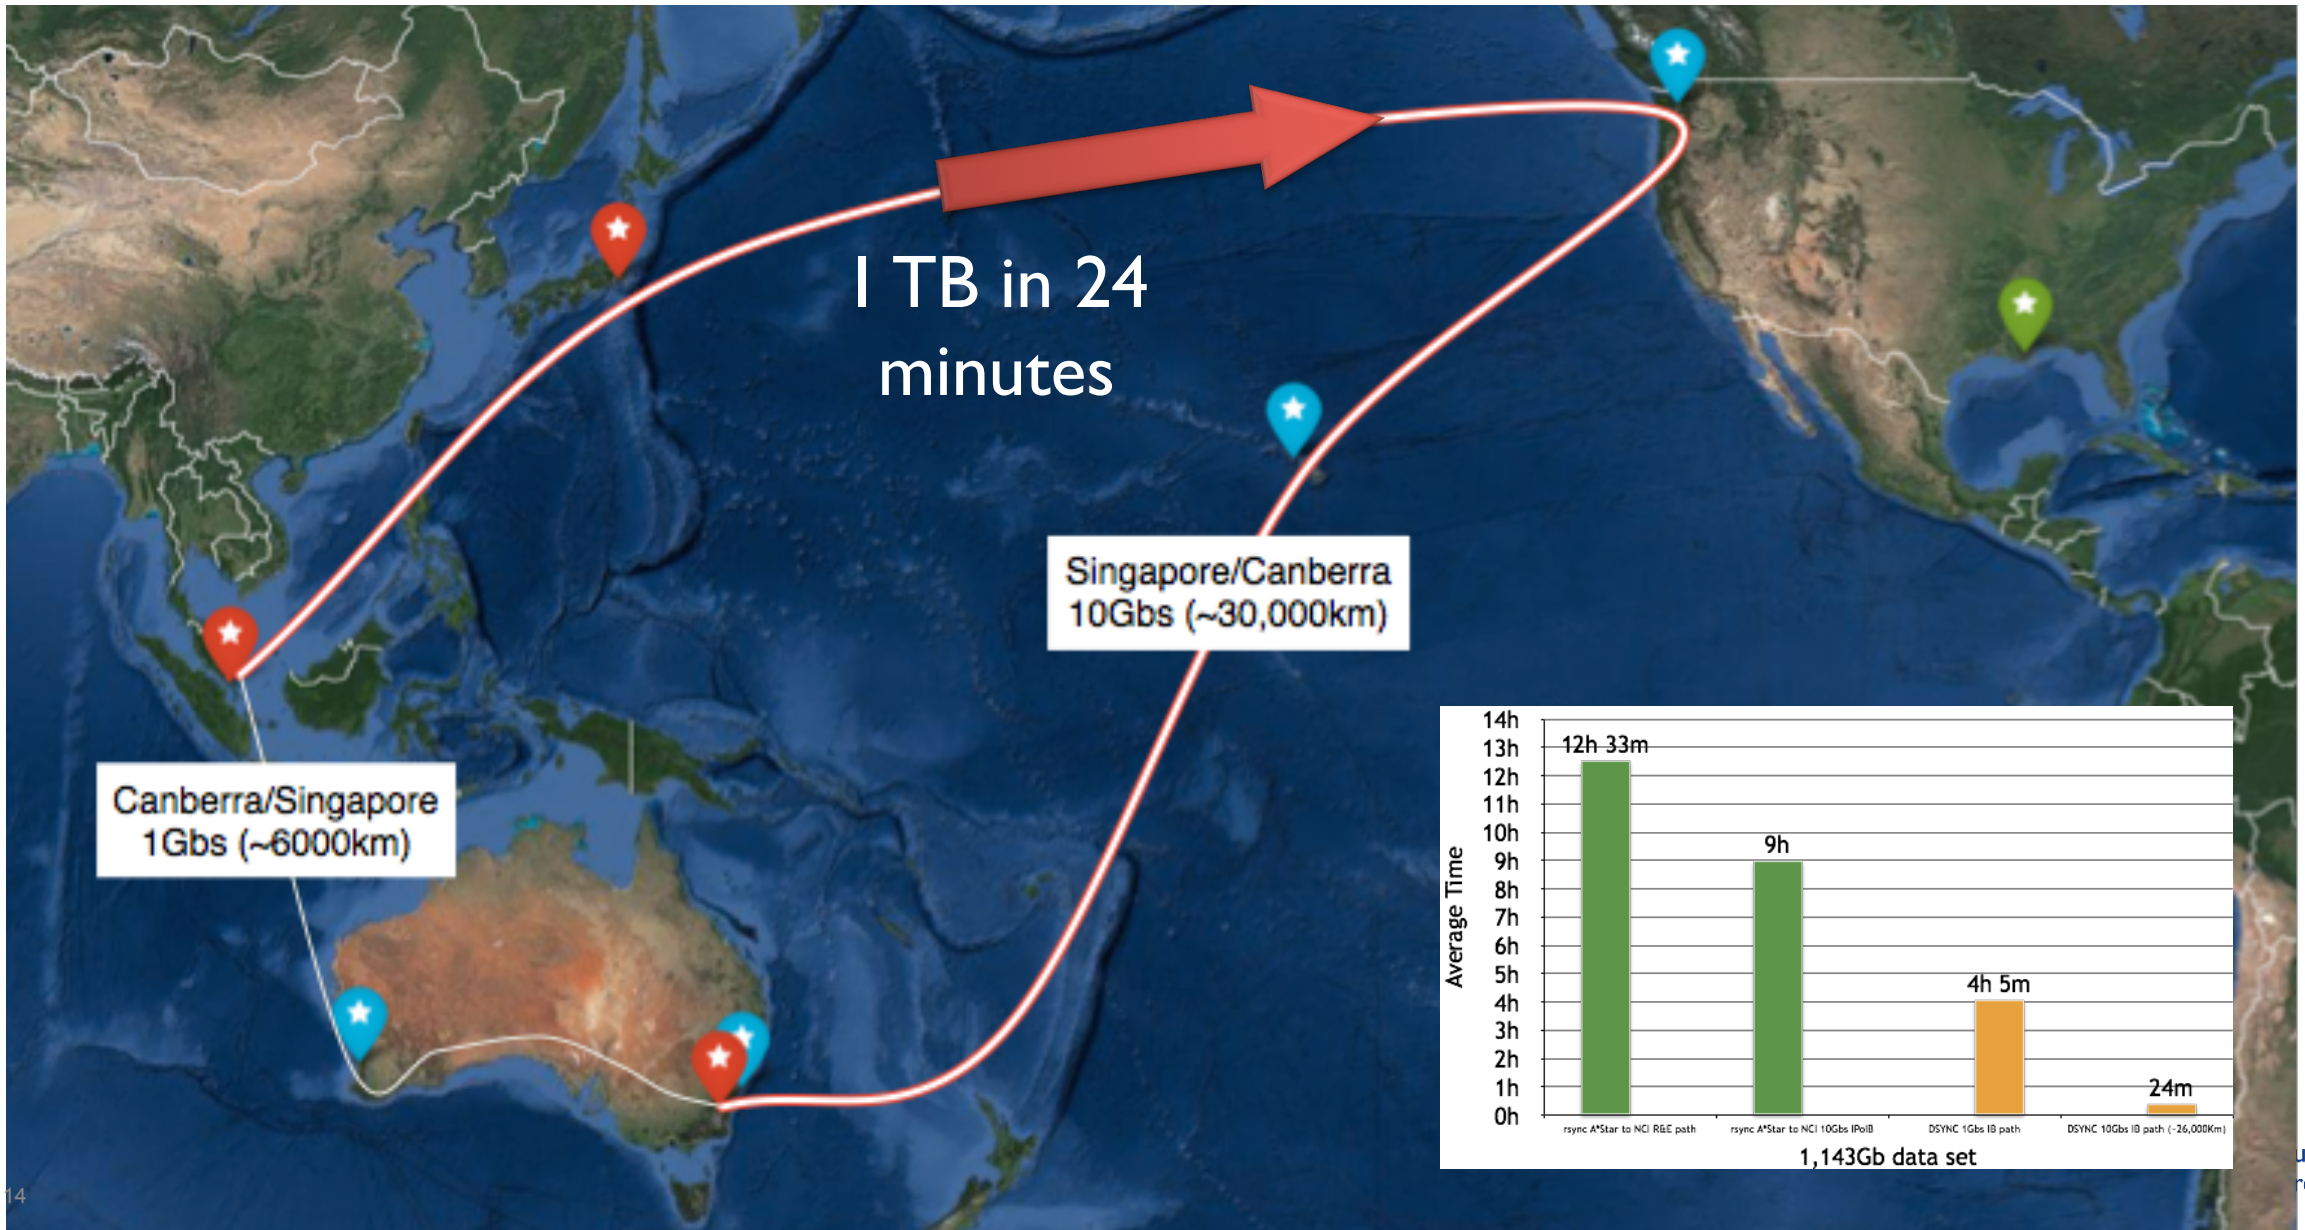
scalable InfiniBand connected HPC cloud instances - InfiniCloud across four countries (Australia, Singapore, USA, France)

InfiniBand routing across 7 subnets over 4 continents

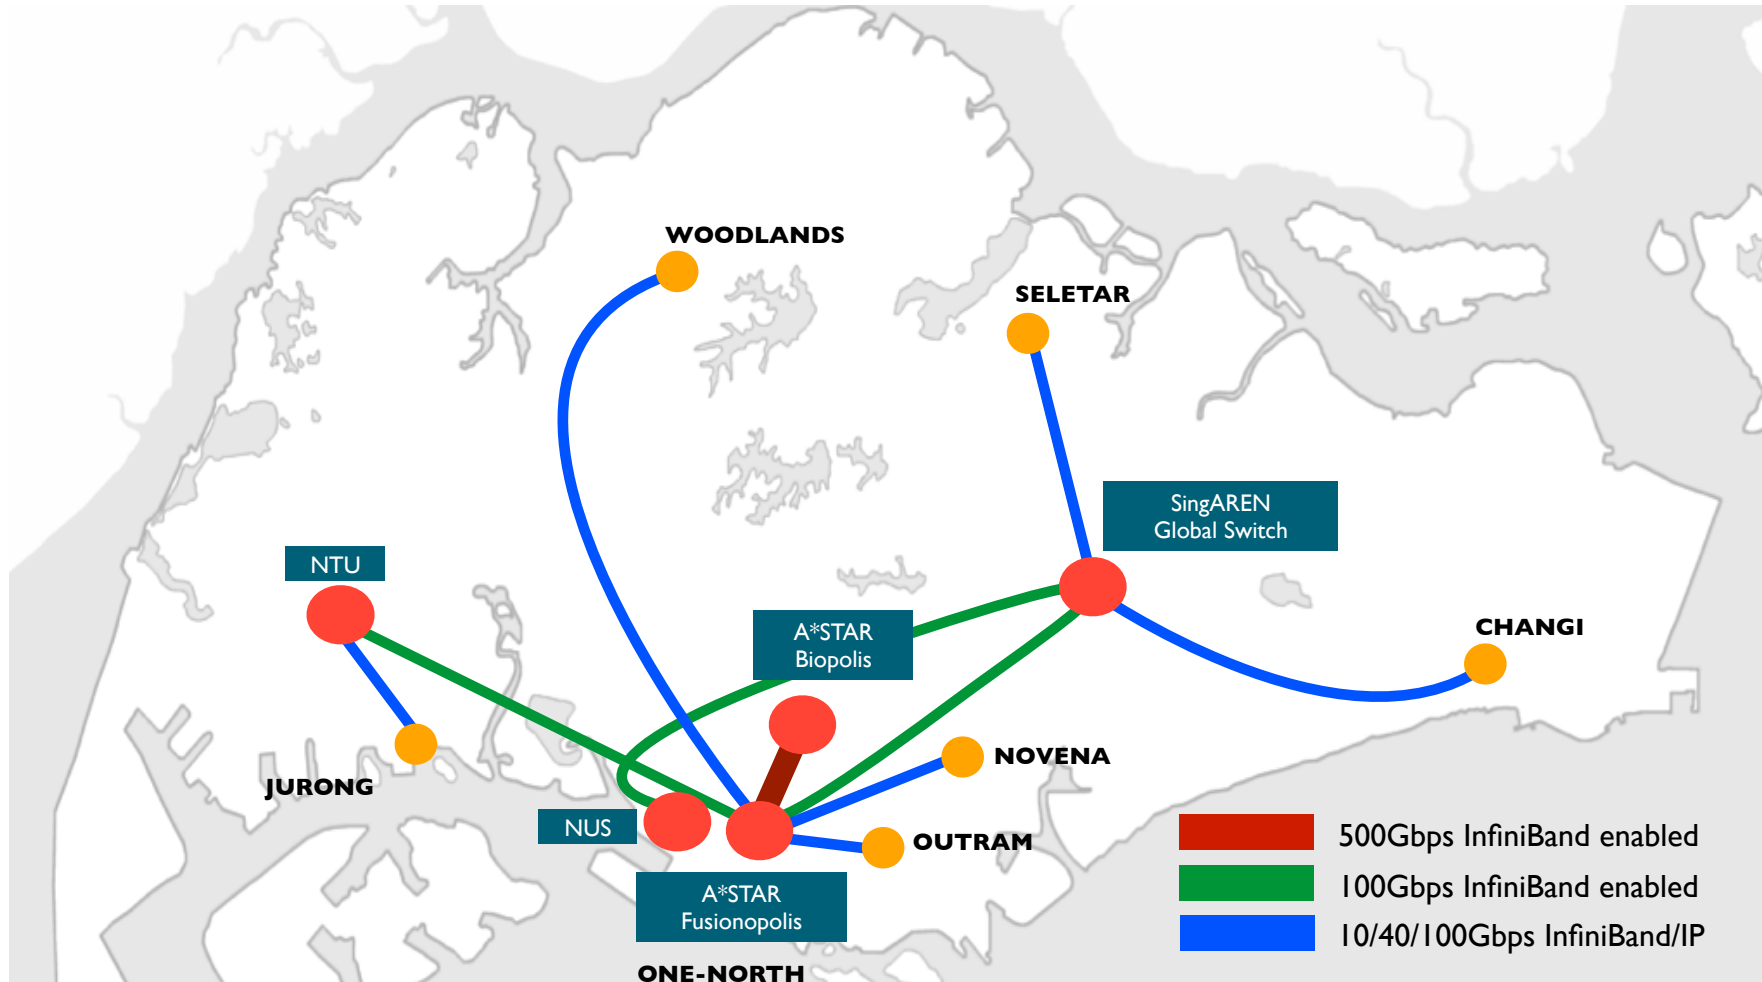
InfiniCloud at ISC 2016 Frankfurt, Germany



Data Transfer Test (dsync)

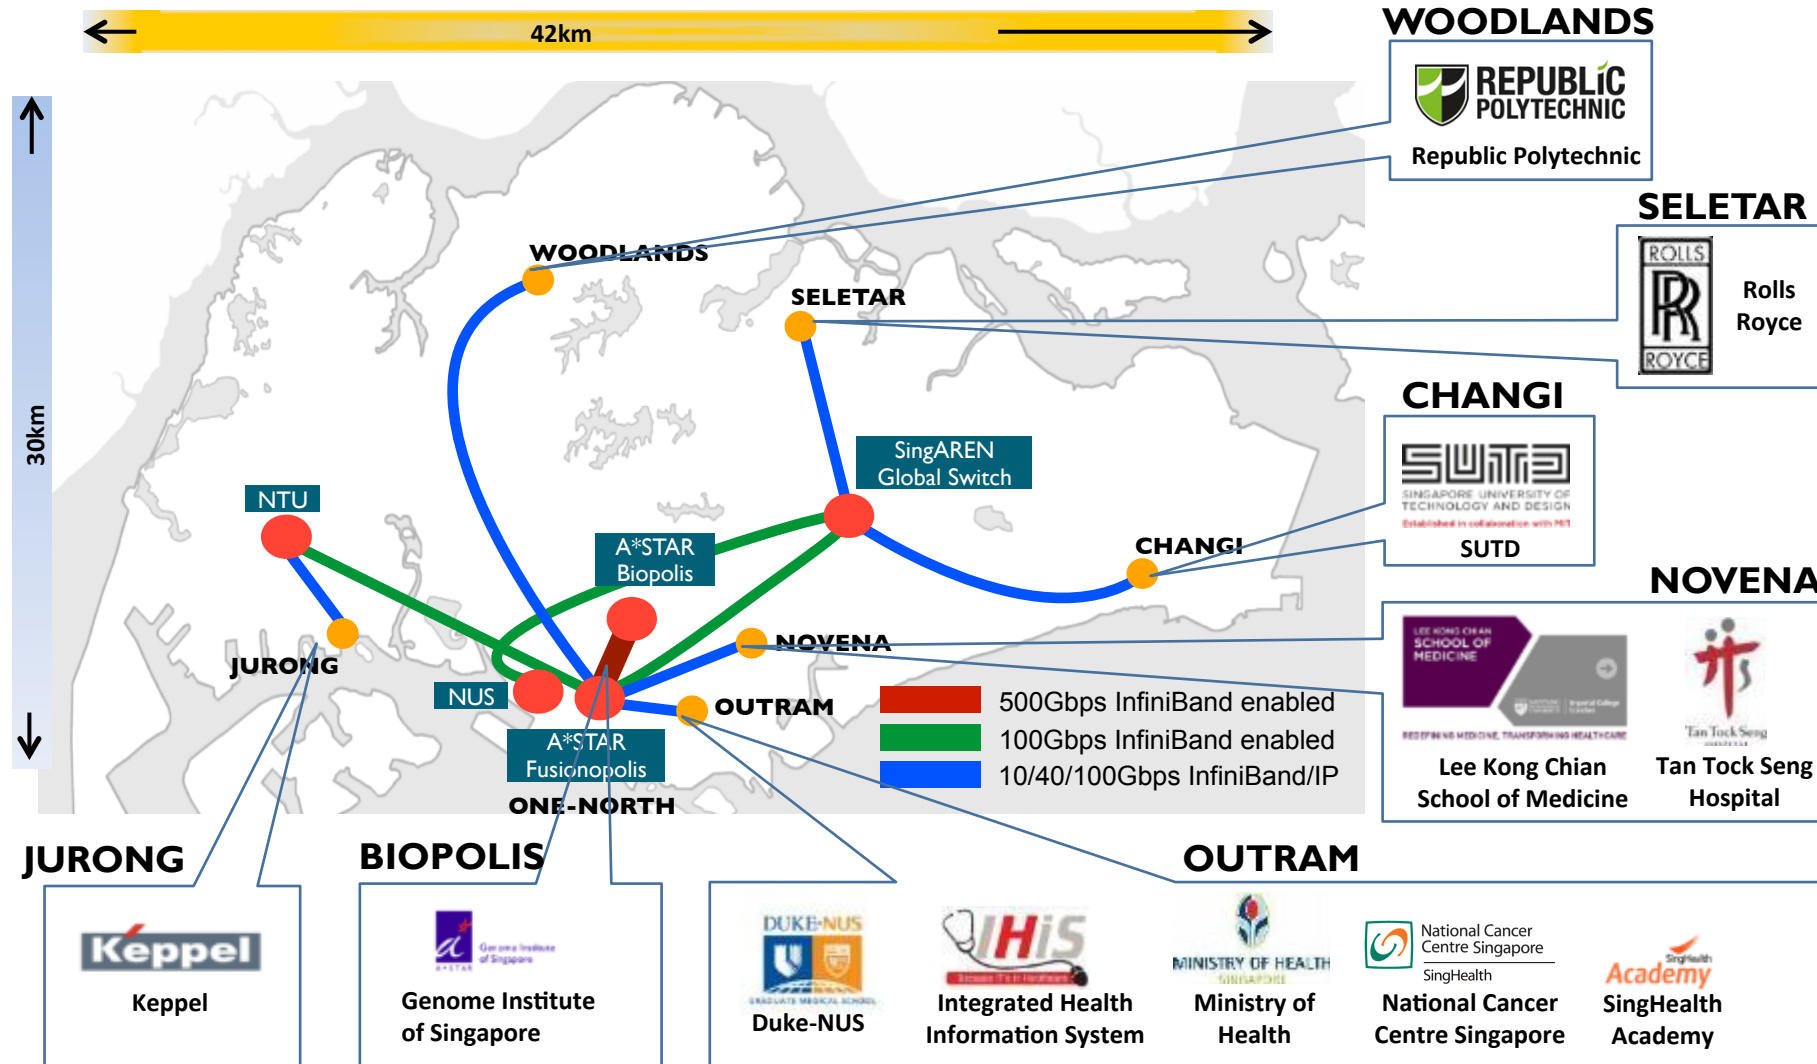


Singapore InfiniBand Connectivity

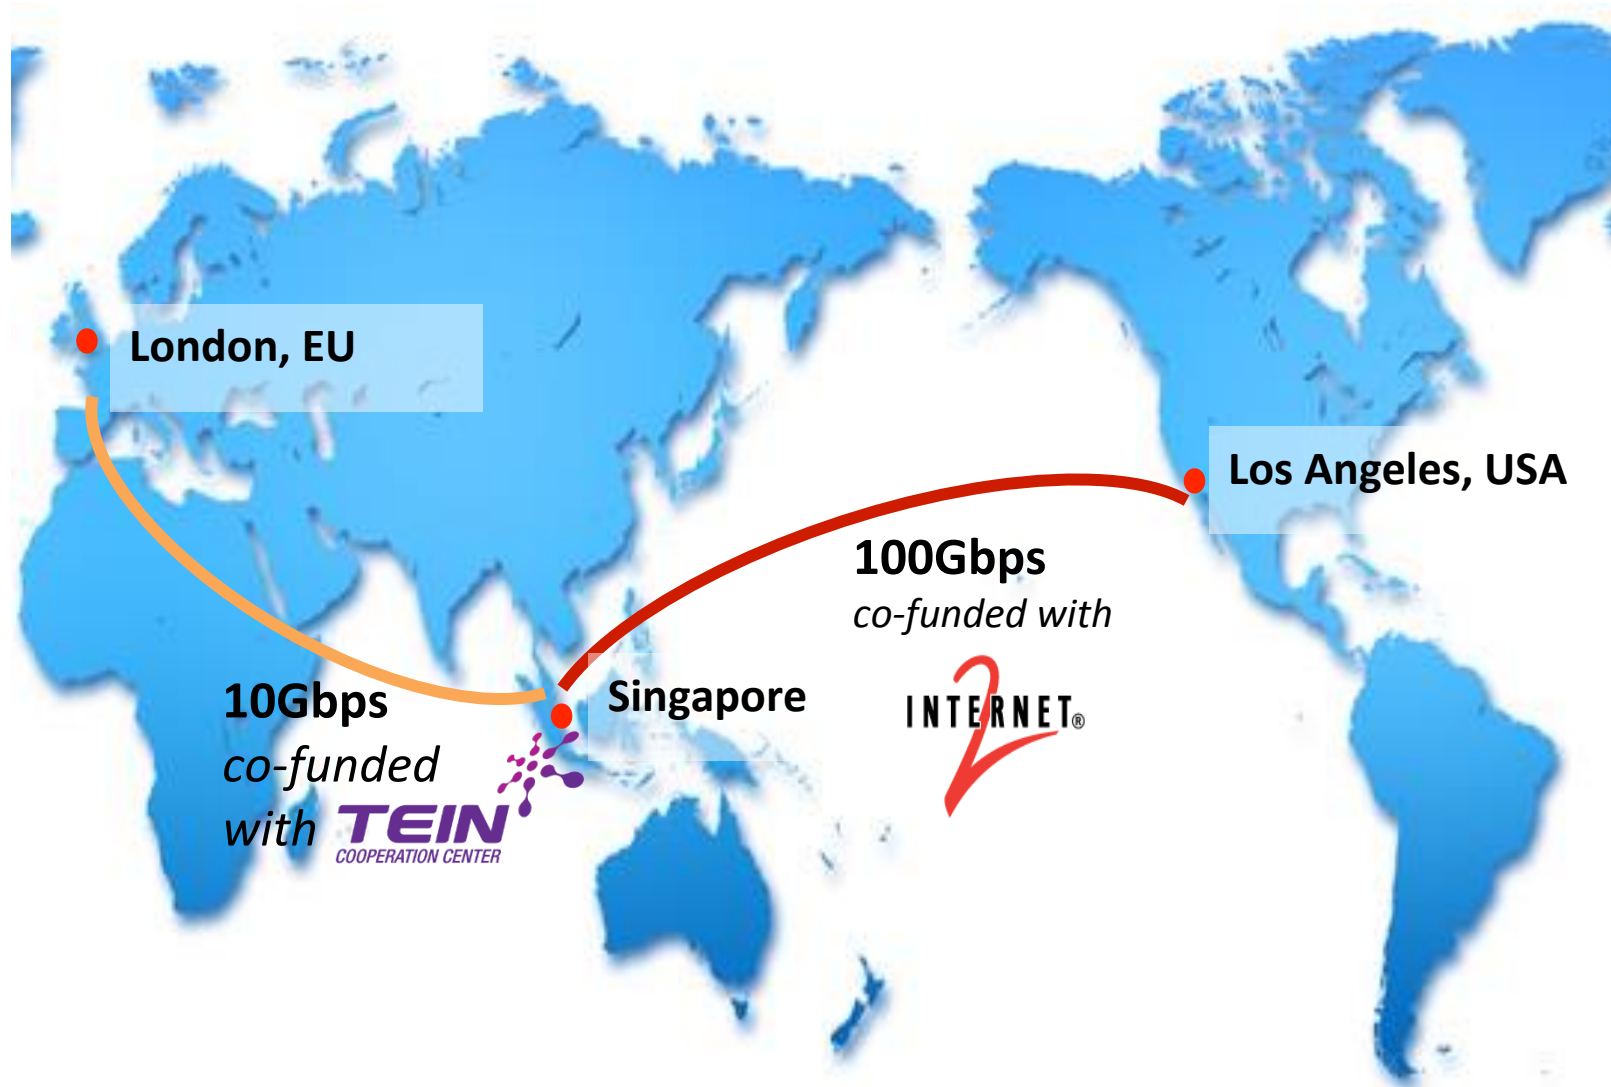


Connecting all National Supercomputing Centre stakeholders:
A*STAR, NUS, NTU, SUTD and others with InfiniBand links.

STAR-N Singapore InfiniBand Fabric



Singapore International Connectivity





Steve Cotter keynote at APAN in Hong Kong:

We are at the confluence of Computing, Storage and Networking

**Creativity requires the courage to let go of certainties.
Erich Fromm**

Thank you!