InfiniCortex

A Path to Reach Exascale Supercomputing

A*STAR Computational Resource Centre

ROMEO 2015

Journee scientifique du Centre de Calcul et de la Maison de la Simulation

June 11, 2015

Dr Marek T. Michalewicz

Dr Gabriel Noaje

Geok Lian Tan

A*STAR Computational Resource Centre

Contents

1. Singapore and A*STAR - an introduction
2. A*STAR Computational Resource Centre
3. Context (Challenges, StarN)
4. InfiniCortex – Vision & Motivation
5. Galaxy of Supercomputers
6. InfiniBand Over Trans-continental Distances
7. Application Layer
8. Initiatives
9. Perspectives

Overview of Singapore

- **Physical Land Area: 714 sq km**
- **Population: 5.18 million**
  - 3.8 million (73%) Singapore Residents
- **Literacy Rate: 96.1%** (Aged 15 & above)
  - 72% aged 25-34 years holds diploma or above
- **2011 GDP S$326.8bn** (US$262bn)
  - Real Growth 4.9%
  - Per Capita GDP S$63,050 (US$50,000)
- **Gained independence on 9 August 1965**

Sources: Singapore in Brief 2012 published by Department of Statistics

Agency for Science, Technology and Research
SINGAPORE

Singapore in Relation to the Asia Pacific Region

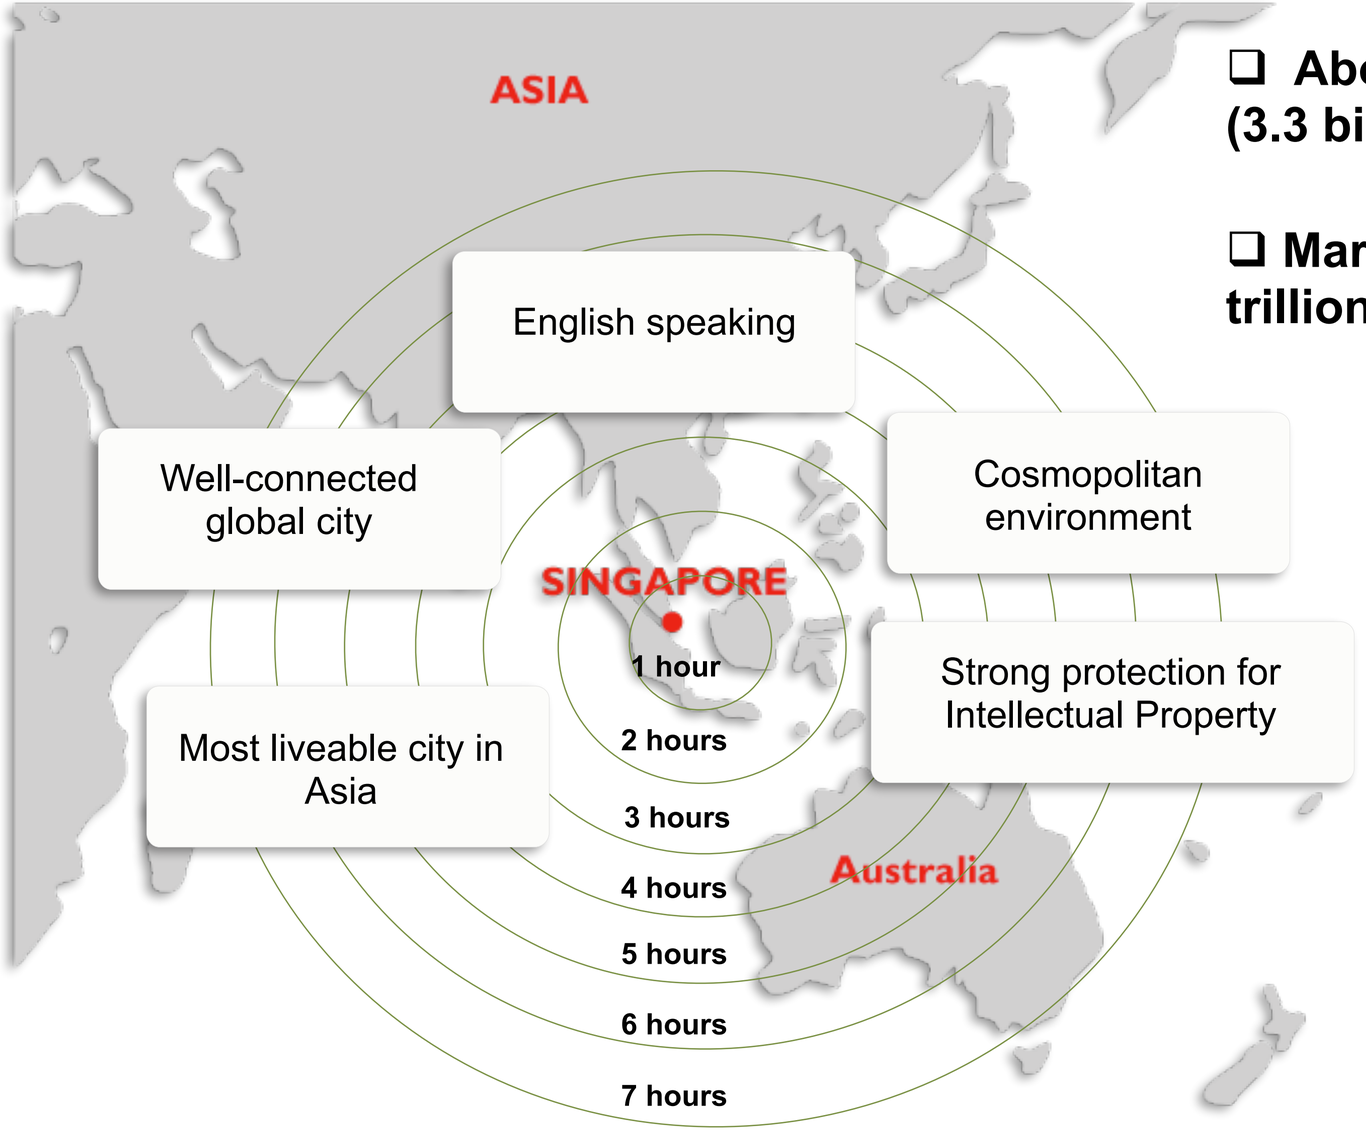

- **About half the world's population (3.3 billion people) live here**
- **Markets worth US$14 trillion (€10 trillion)**

The Singapore Economy

Manufacturing accounts for 20.9% of GDP and grew 7.6% in 2011

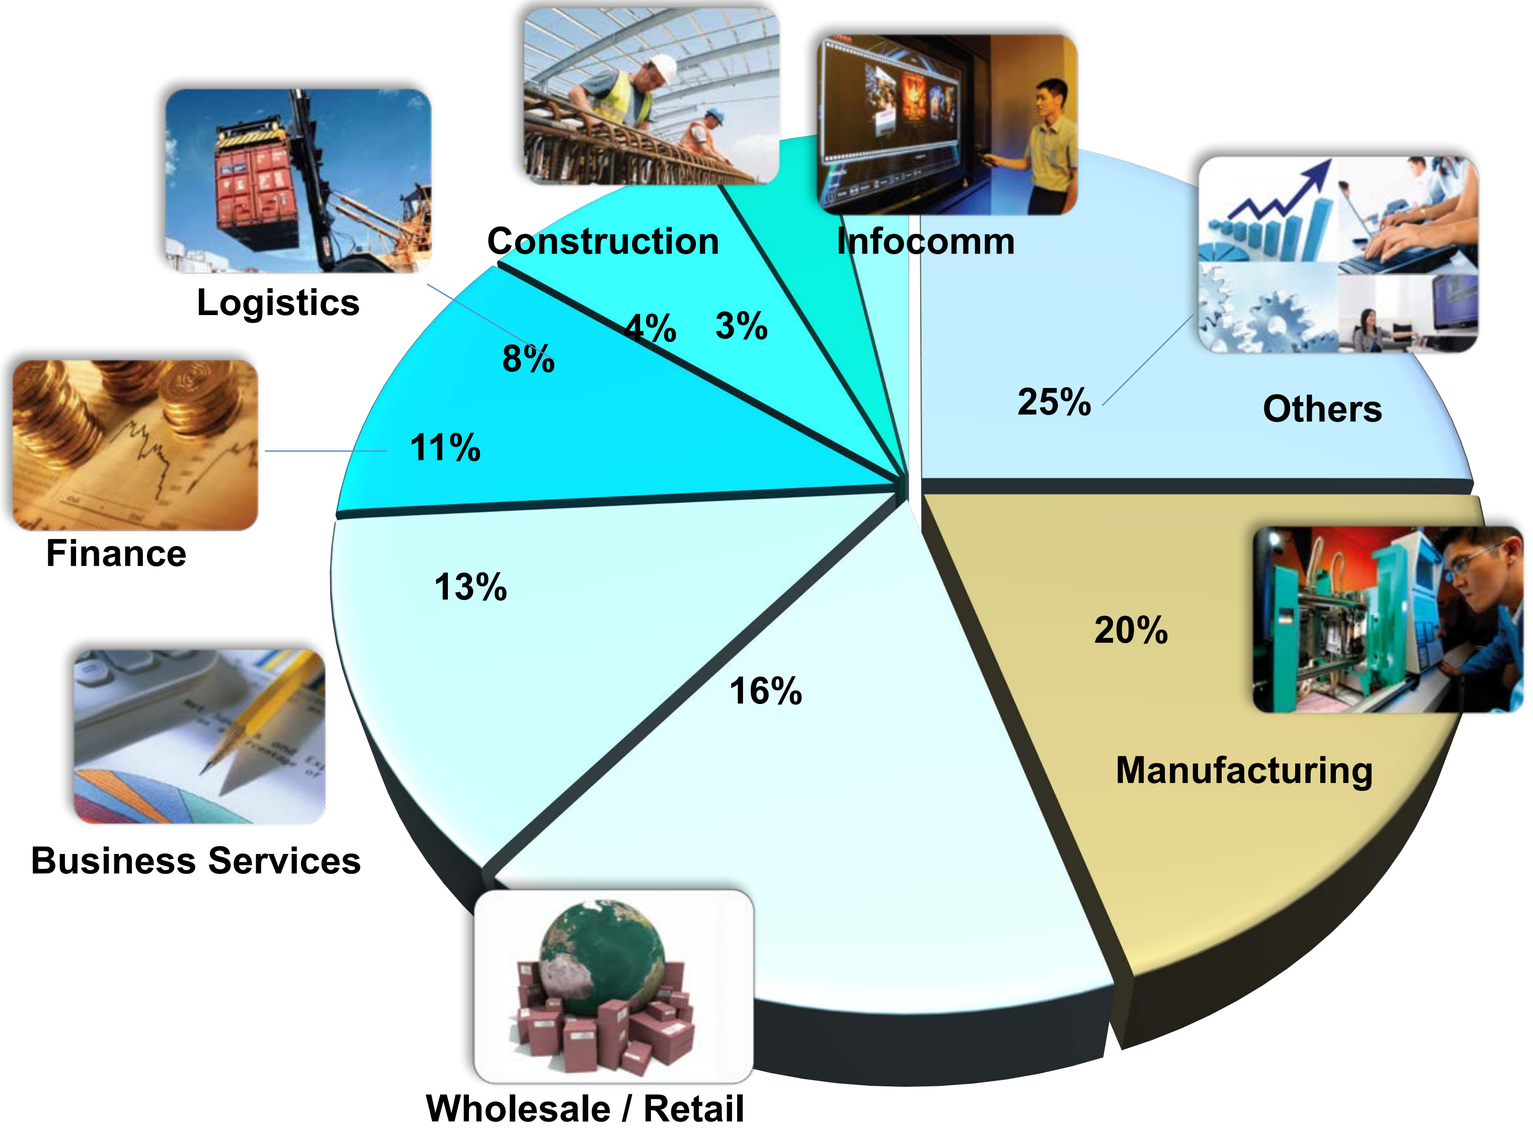

Agency for Science, Technology and Research
SINGAPORE

Growth of the Singapore Economy

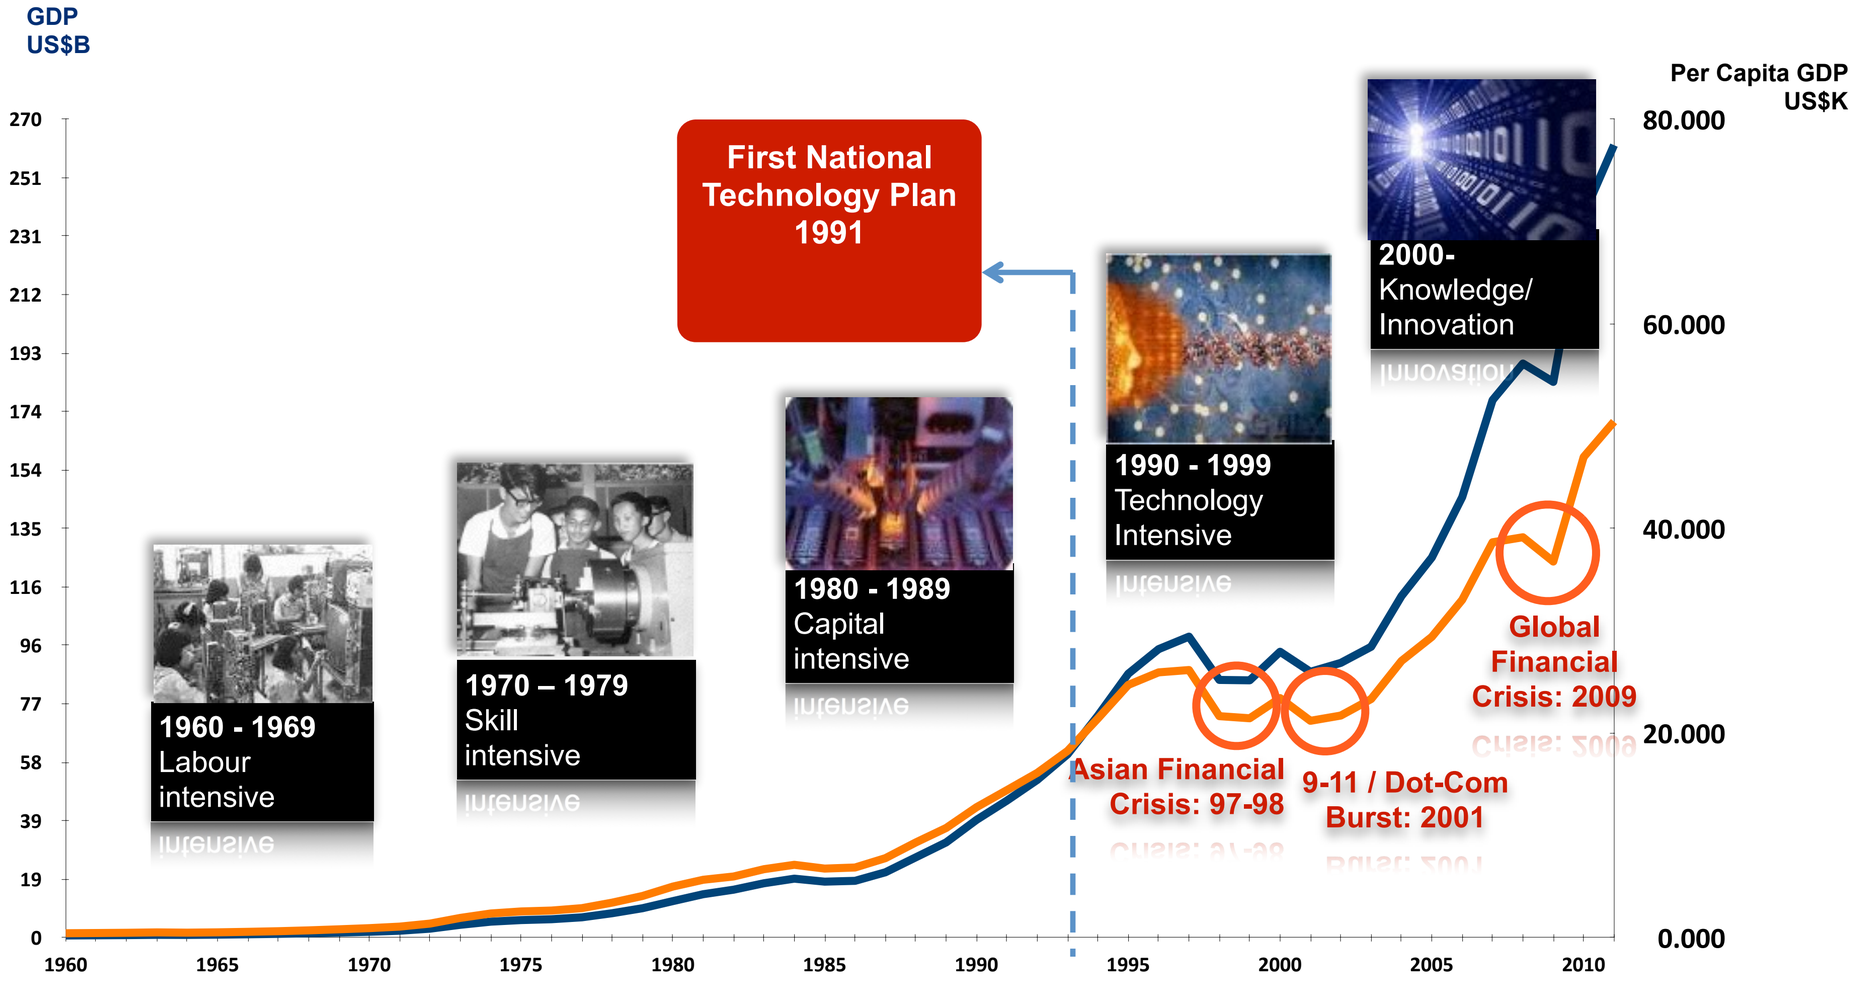

GDP
1960 US$704.5 mil
2011 US$262 bn
(S$326.8 bn)

Per Capita GDP
1960 US$428
2011 US$50,000 (S $63,050)

Source: Singapore Ministry of Trade and Industry

Agency for Science, Technology and Research
SINGAPORE

R&D Expenditure in Singapore

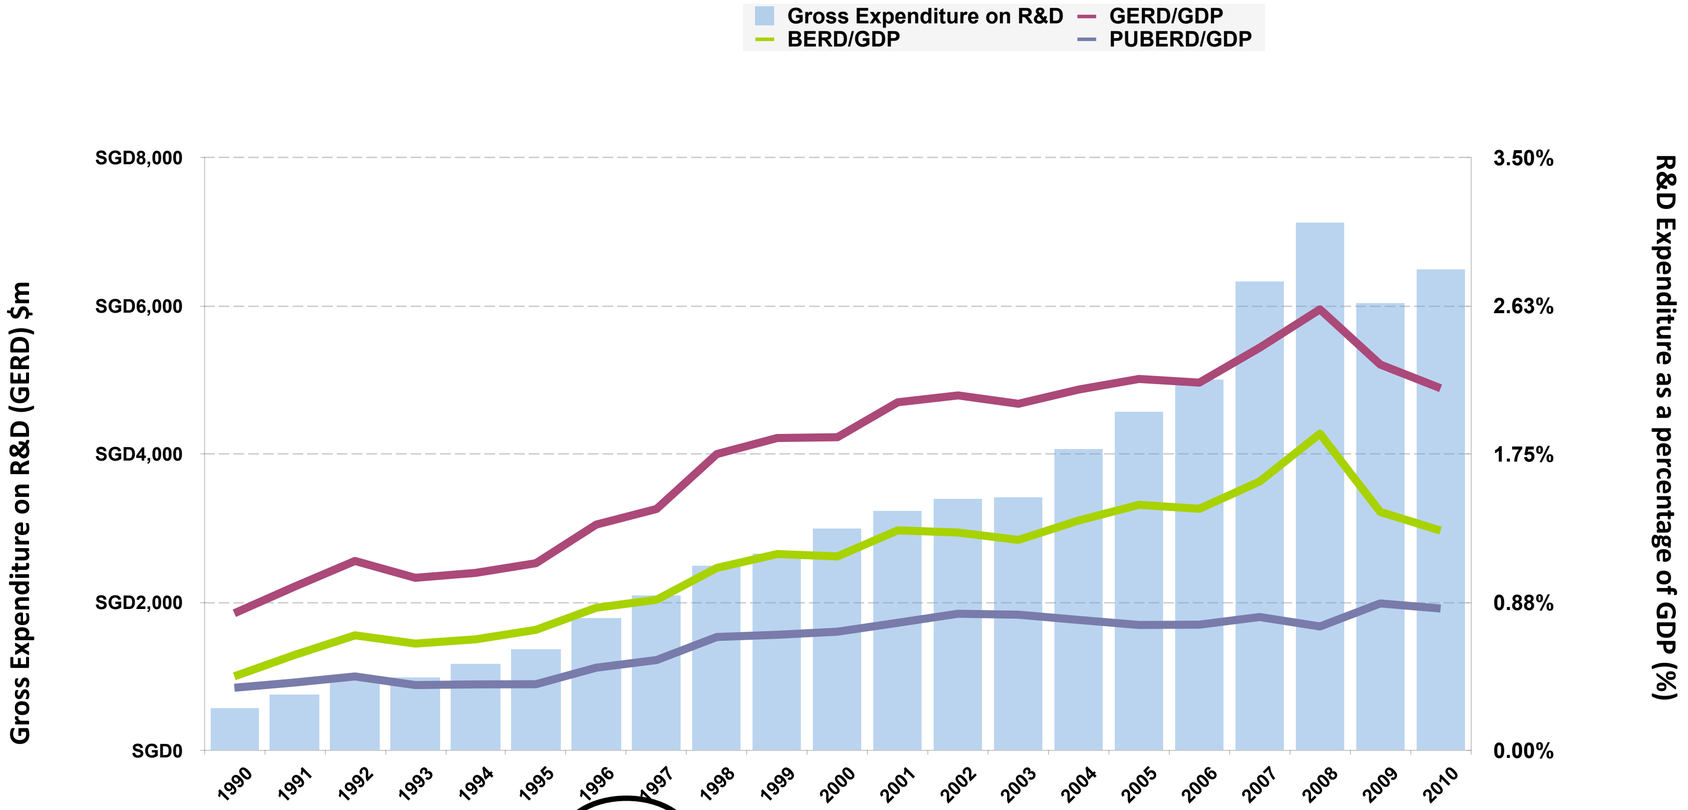

"Singapore aspires to achieve 3.5% GERD/GDP, this will place us amongst the most research intensive countries in the world"

Agency for Science, Technology and Research
SINGAPORE

Agency for Science, Technology and Research
SINGAPORE

Our Mission

Fostering world-class scientific research and talent for a vibrant knowledge-based Singapore

Chairman A*STAR
Managing Director A*STAR

SERC

BMRC

A*STAR Computational Resource Centre

Data Storage Institute (DSI)
Institute of Chemical and Engineering Sciences (ICES)
Institute of High Performance Computing (IHPC)
Institute for Infocomm Research (I^2R)
Institute of Materials Research and Engineering (IMRE)
Institute of Microelectronics (IME)
National Metrology Centre (NMC)
Singapore Institute of Manufacturing Technology (SIMTech)

Bioinformatics Institute (BII)
Bioprocessing Technology Institute (BTI)
Genome Institute of Singapore (GIS)
Institute of Bioengineering and Nanotechnology (IBN)
Institute for Medical Biology (IMB)
Institute of Molecular and Cell Biology (IMCB)
A*STAR-Duke-NUS Graduate Medical School Neuroscience Research Partnership (NRP)
Singapore Bioimaging Consortium (SBIC)
Singapore Consortium of Cohort Studies (SCCS)
Singapore Institute for Clinical Sciences (SICS)
Singapore Immunology Network (SIgN)
Singapore Stem Cell Consortium (SSCC)
Experimental Therapeutics Centre (ETC)

Commercialisation Outfits

- Exploit Technologies
- Experimental Therapeutics Centre
- Industry Development Group (BMRC)

Other Scientific Services Outfits

- Biological Resource Centre
- National Breeding Centre
- Singapore Tissue Network
- National Metrology Centre
- Biopolis Shared Facilities

Reasons for A*CRC's Existence

1. Service to A*STAR HPC users

2. Service to A*STAR HPC users

3. Service to A*STAR HPC users

A*CRC is not actively involved in own research, nor product development, but:

- we study state-of-the-art HPC technologies,
- we engage in forward trends discussions with vendors
- we observe and study best practices and trends
- we implement the best technological solutions that suit our users' needs.

A*STAR Computational Resource Centre

Users Statistics

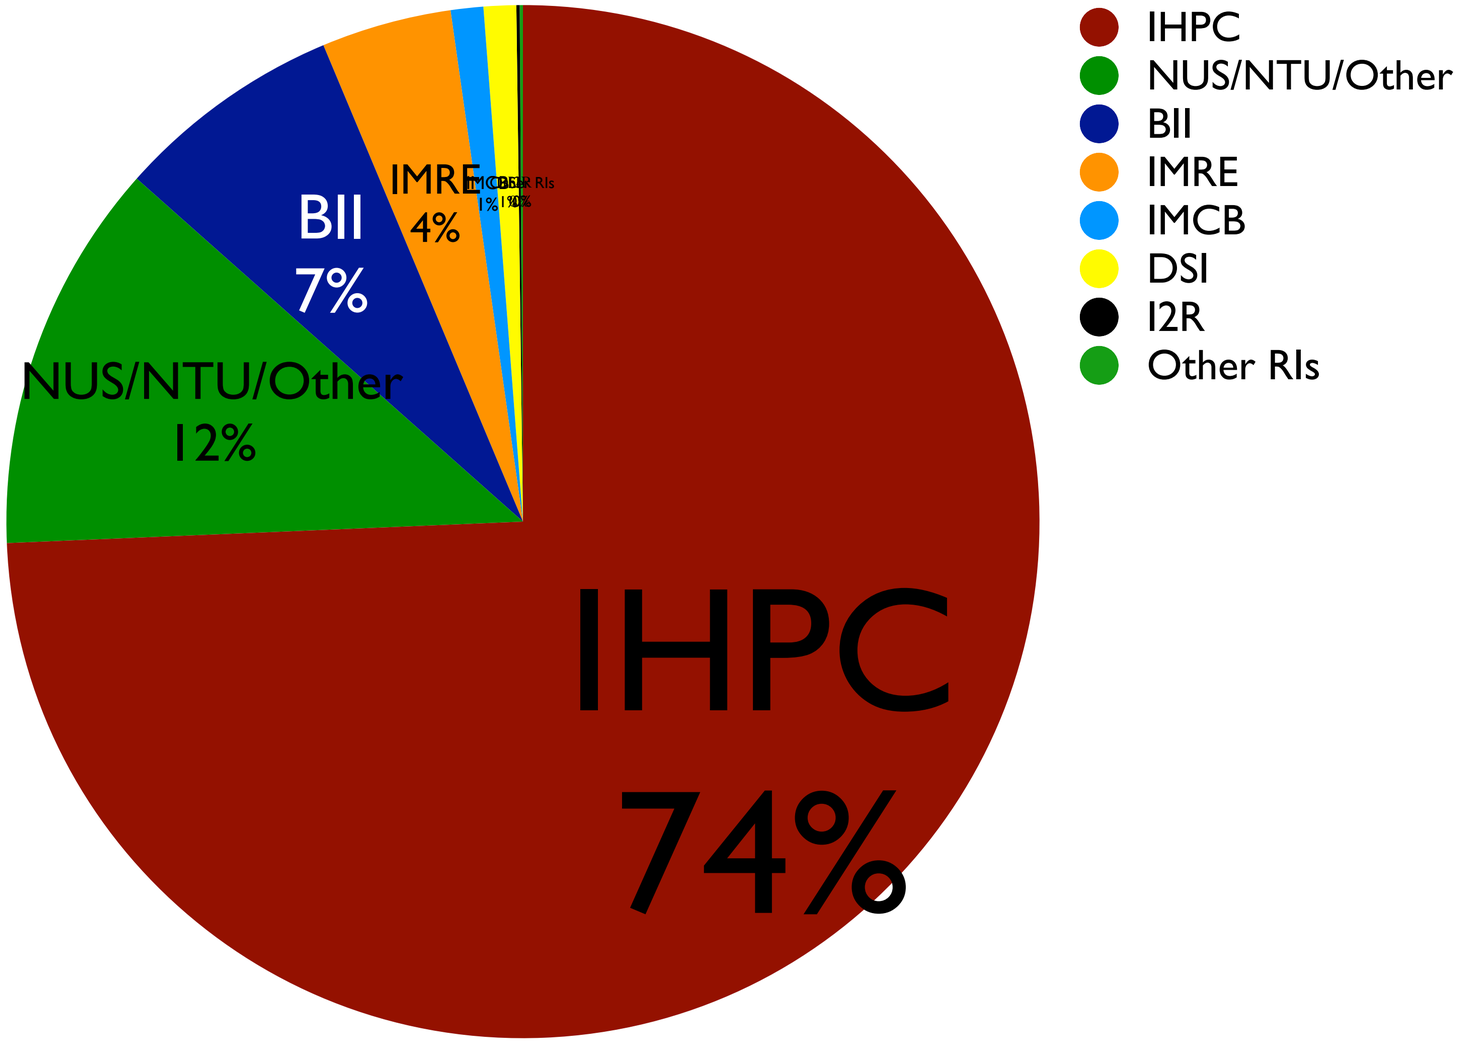

- **~800 users**
- **5300+ A*STAR staff**
- **all A*STAR Institutes and Centres,**

  but
- **largest group from the Institute Of High Performance**

  Computing

A*CRC Compute Systems

SYSTEM NAME	Cumulus	Aurora (A)	Aurora (B)	Fuji	Axle	Cirrus
AGE [years]	1	3	3	4	5	2.5
LIFE LEFT [years]	3	1	1	0.5	0	1.5
VENDOR	IBM	SGI	SGI	Fujitsu	HP	IBM
SYSTEM ARCHITECTURE	Linux	SMP	SMP	cluster	cluster	cluster
OPERATING SYSTEM	Linux	Linux	Linux	Linux	Linux	AIX
NUMBER OF NODES	512	1	1	450	32	30
CORES PER NODE	16	2,048	576	8	32	32
TOTAL CPU CORES	8192	2,048	576	3,888	1,024	960
MEMORY/NODE [GB]	16	12,000	3,000	24	128	128
TOTAL MEMORY [GB]	8192	12,000	3,000	9,331	4,096	3,840

A*STAR HPC Computational Power

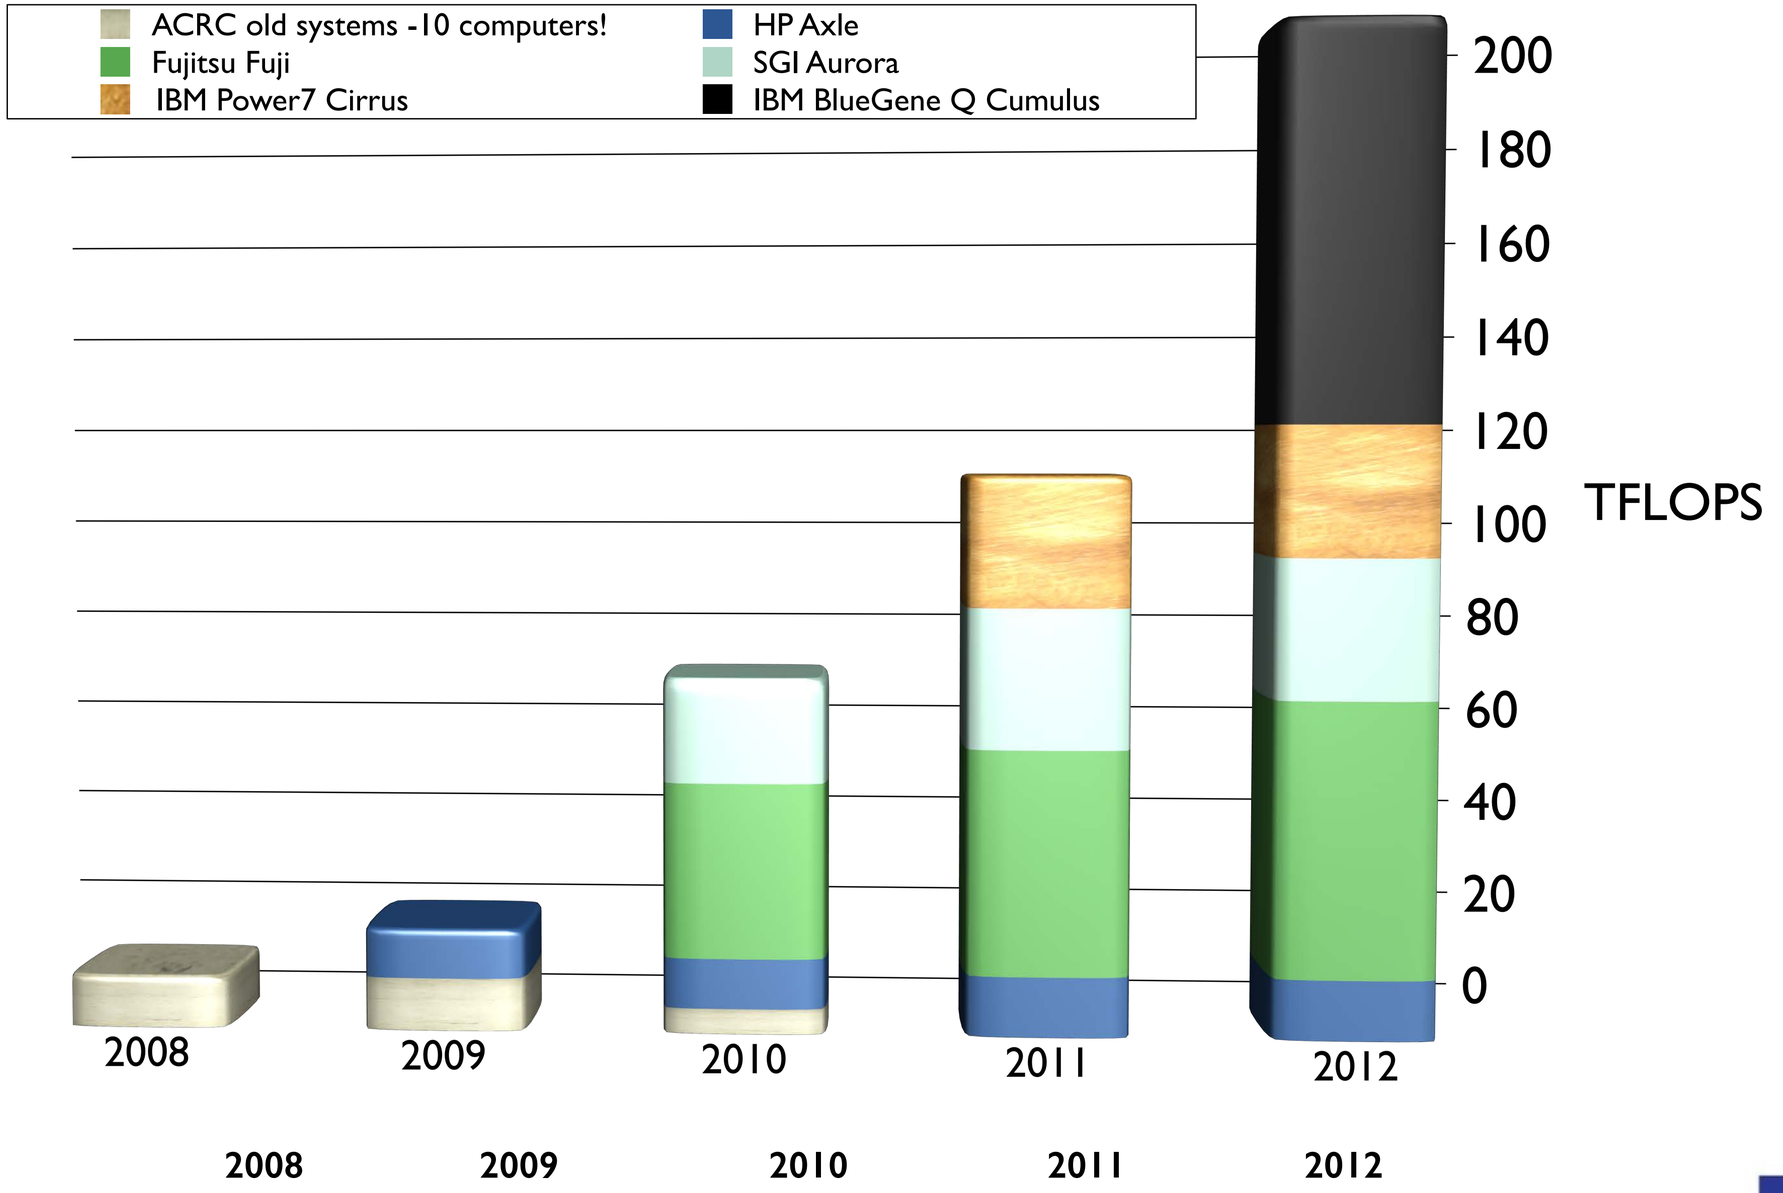

Agency for Science, Technology and Research
SINGAPORE

A*CRC Organisation Structure

Prof Tan Tin Wee
Chairman

Dr John Kan
Deputy Chairman

Kekwa Bte
Mohammed
Secretary

Dr. Marek T Michalewicz
Senior Director

Software

Dr Chien Siu-Hung
Software Manager

Liou Sing-Wu
Visiting Scientist

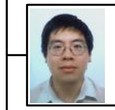

Dr Jonathan Low
Senior Computational
Scientist

Dr Gabriel Noaje
Senior Computational
Scientist

Data Centre & Network

Tan Geok Lian
Network Manager

Lim Seng
Senior Network
Engineer

New Staff

Storage & Cloud Services

Lai Loong Fong
Director

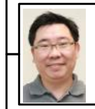

Hoi Wai Kok, Kenny
Deputy Director

Lin Yi, Dave
Systems Manager

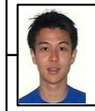

Yap Chung Lam
Systems Manager

Yeo Chang Lim
Systems Manager

Chew Shiun Han
Asst Systems Manager

Operations

Lim Ching Kwang
Deputy Director

Ho Wai Sing, Edmund
Assistant Systems
Manager

Mohd Jamaludeen
Systems Engineer

Computing Systems

Hiew Ngee Heng, Paul
Systems Manager

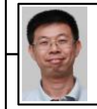

Kenneth Koo Guan Sim
Systems Manager

Tay Teck Wee
Systems Manager

Steven Chew Beng Hwee
Senior Systems Engineer

Kevin Siswandi
Systems Engineer

Yee Lee Poh
System Engineer

A*STAR Computational Resource Centre

A*CRC Datacenter 1

A*CRC Datacenter 2

Metro-X A*CRC team:
Stephen Wong
Tay Teck Wee
Steven Chew

Mellanox Metro-X testing since early 2013

goal: to connect HPC resources at Fusionopolis with storage and genomics pipeline at Biopolis - Matrix building

A*STAR Computational Resource Centre

Early tests with MetroX Mellanox switches

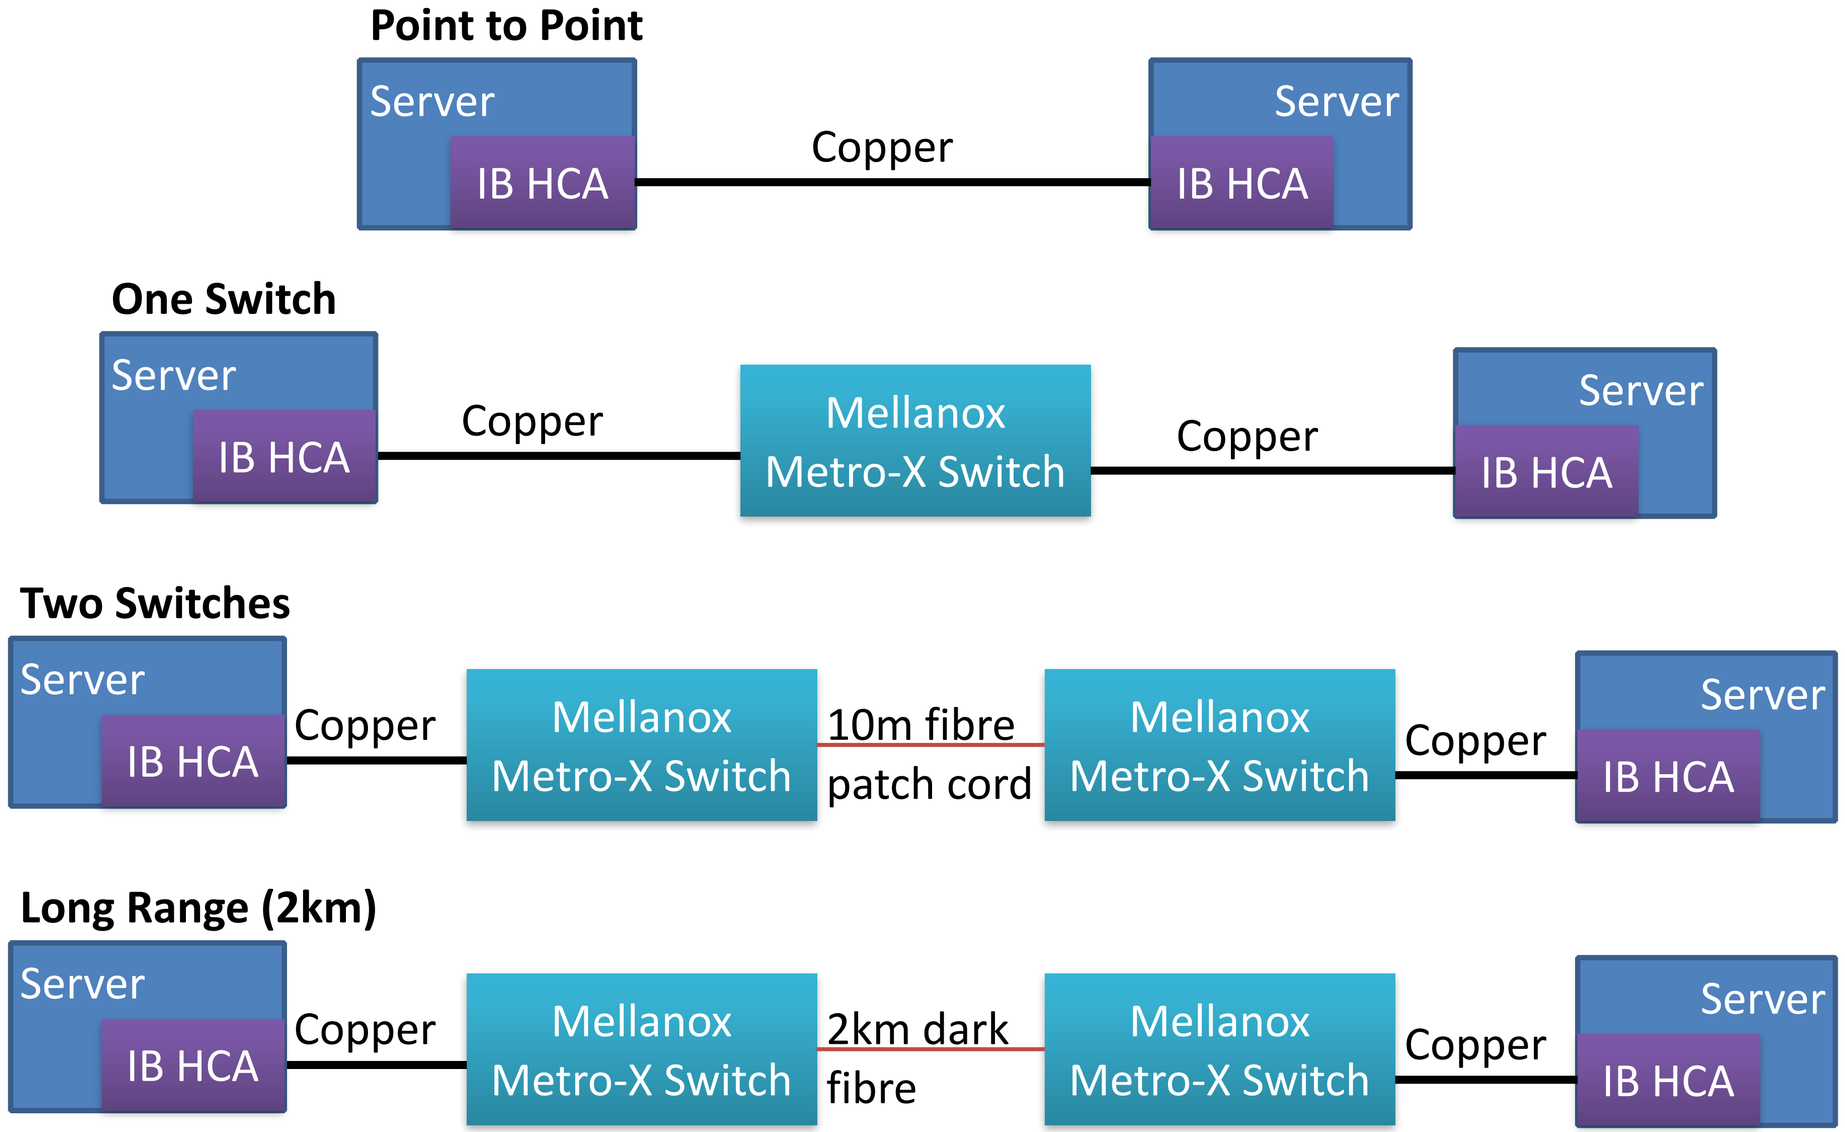

National Supercomputer Centre

Joint A*STAR, NUS, NTU, SUTD and NRF Proposal

- National Supercomputing Centre (NSCC)
  - New 1-3+ PetaFLOP Supercomputer
  - Recurrent investment every 3-5 years
  - Co-investment from primary stakeholders
- Science, Technology and Research Network (STAR-N)
  - A high bandwidth network to connect the distributed compute resources
  - Provide high speed access to users (both public and private) anywhere
  - Support transfer of large data-sets (both locally and internationally)
  - Build local and international network connectivity

Funding (MTI) and co-funding (A*STAR, NUS, NTU) approved Nov. 2014
Tender open 20th January 2015
Tender Closed 14th April
Facility open to users: October 2015

A*STAR Computational Resource Centre

Some features of NSCC Supercomputer

1. Base system: ~1 PFLOPS
2. Storage ~10PB, HSM, Tiered, Tier 3 - tape storage (or very cheap, non-spinning disks)
3. I/O comparable to the best systems around (BlueWaters class)
4. Application Software - National license desirable - ISVs, MATLAB, tools e.g. Allinea, NAG
5. 10 Large memory nodes: 5 nodes of 1TBytes, 4 nodes with 2-4TBytes and one node of 6-8TBytes
6. Workflow pipeline (from sequencer to memory/storage) + interactive access built in
7. **500 Gbps pipeline** between Biopolis and Fusionopolis for genomics workflows
8. InfiniBand connection to all end-points (login nodes) at University campuses
9. Just-in-time resource delivery i.e. ***interactive access*** for some workloads (genomics)
10. Only ~10% nodes GPU accelerated (for some University users, AI work - deep learning)
11. Warm water cooled

A*STAR Computational Resource Centre

GIS-NSCC Integration: Future

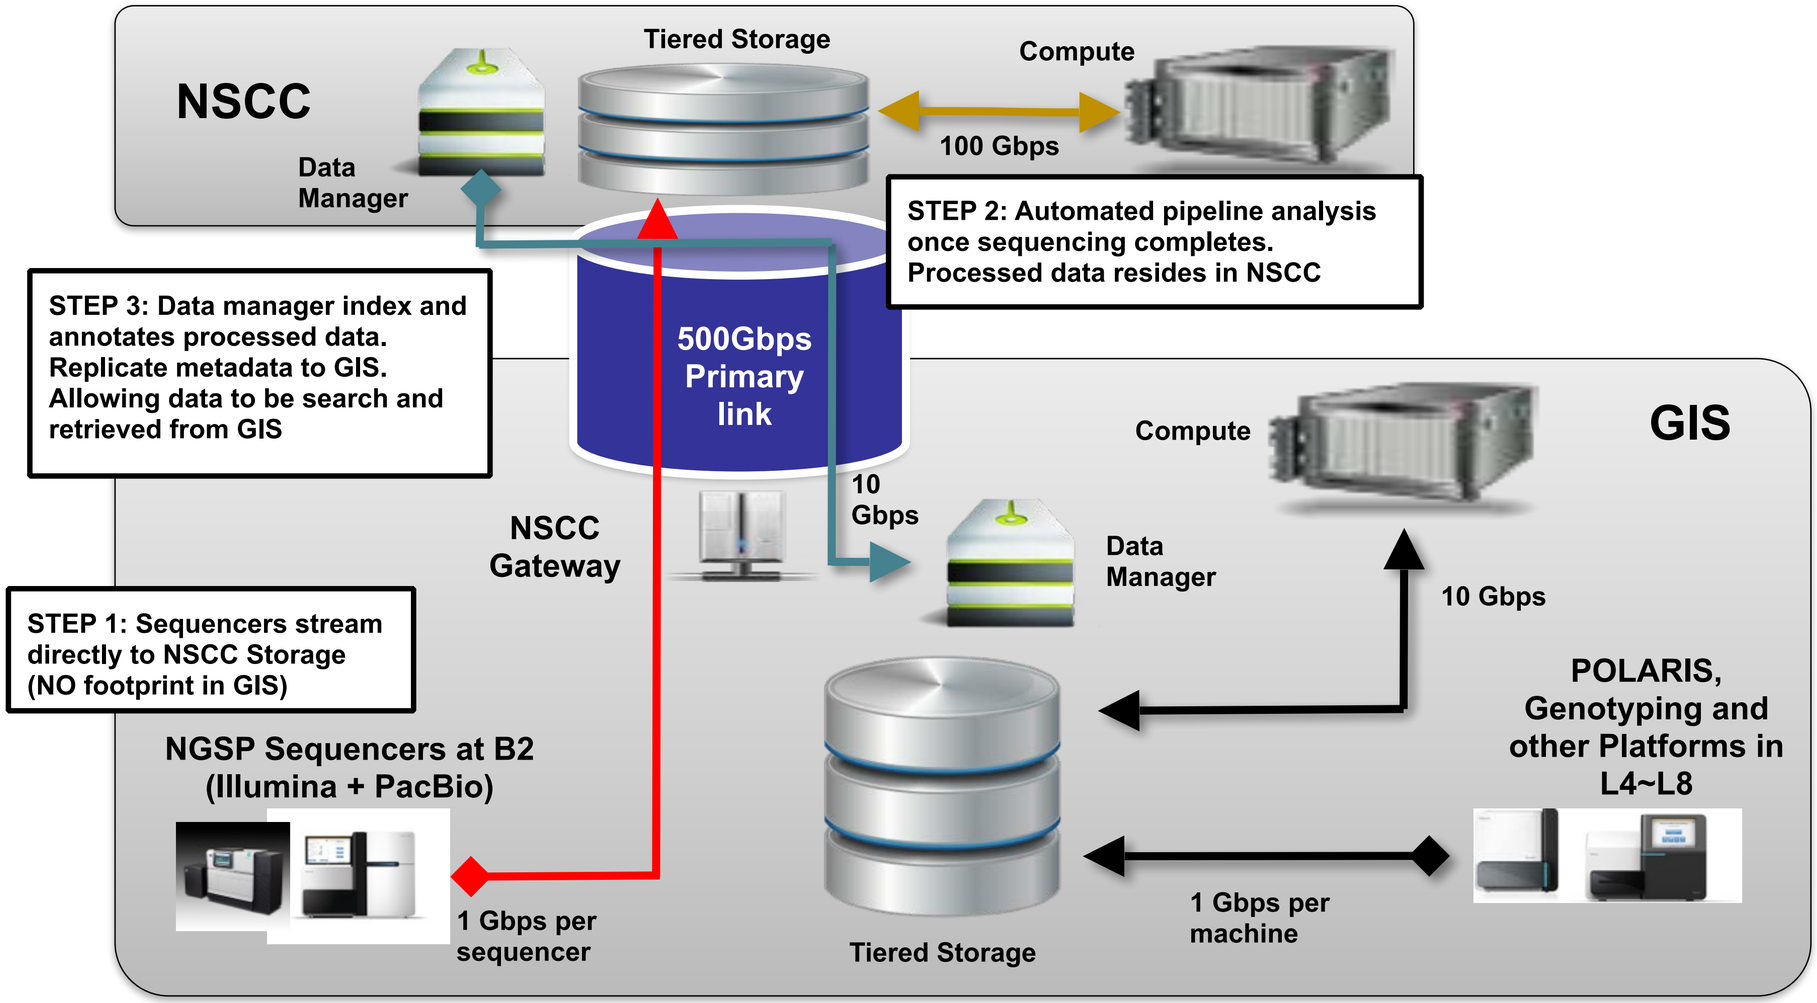

Raising Asian Research and Education Networking to a higher dimension

ACA-100 challenge

Asia connects America at 100Gbps in November 2014 at SC14 in New Orleans, USA

Credits: Yves Poppe, APAN 37 at Bandung, Indonesia January 20, 2014

First US Long Distance InfiniBand Connection

Capability: Bulk, Secure Data Migration at the File System Level

- NASA helped Obsidian define market requirements for enhancing Wide Area InfiniBand connections with suitable encryption and authentication.
- Obsidian created a large-scale file system synchronization tool, dsync+, to help NASA and others simplify the transfer of huge scientific data sets across distance using Longbows – supporting non-InfiniBand storage arrays providing they are fast enough to keep up.

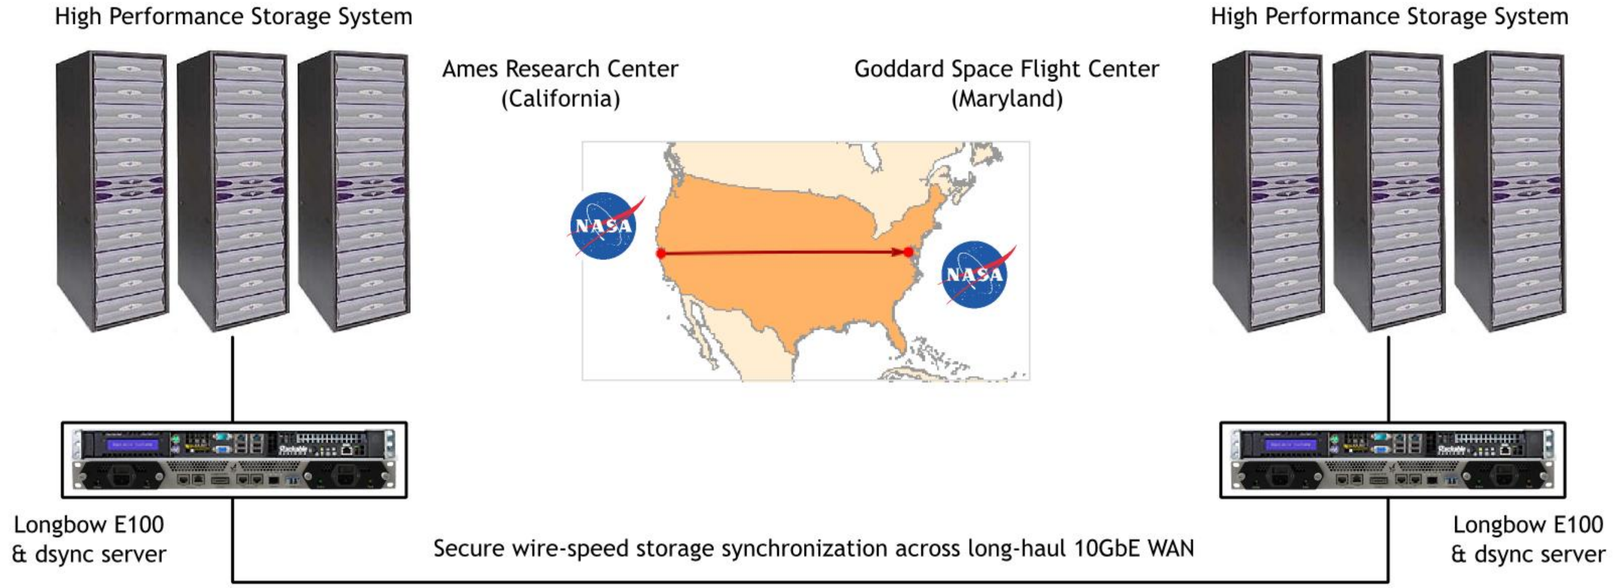

Unprotected ftp transfers – **30 Mbytes/second** file-level copies.

Secure dsync+/Longbow E100 transfers – **940 Mbytes/second** file-level copies.

OBSIDIAN STRATEGICS

A*STAR Computational Resource Centre

NASA Proof of Concept

Capability: Low Latency Server Aggregation

- NASA Ames purchased 16 Longbow C100 units expanding their flagship Itanium-based *Columbia* supercomputer to share jobs across one-mile of dark fiber to a second building.
- Expansion of supercomputers and data centers must contend with power and cooling constraints – these problems can often be resolved by Longbows.
- A similar model works for the linking of containerized data center pods in the field or within modular data centers.

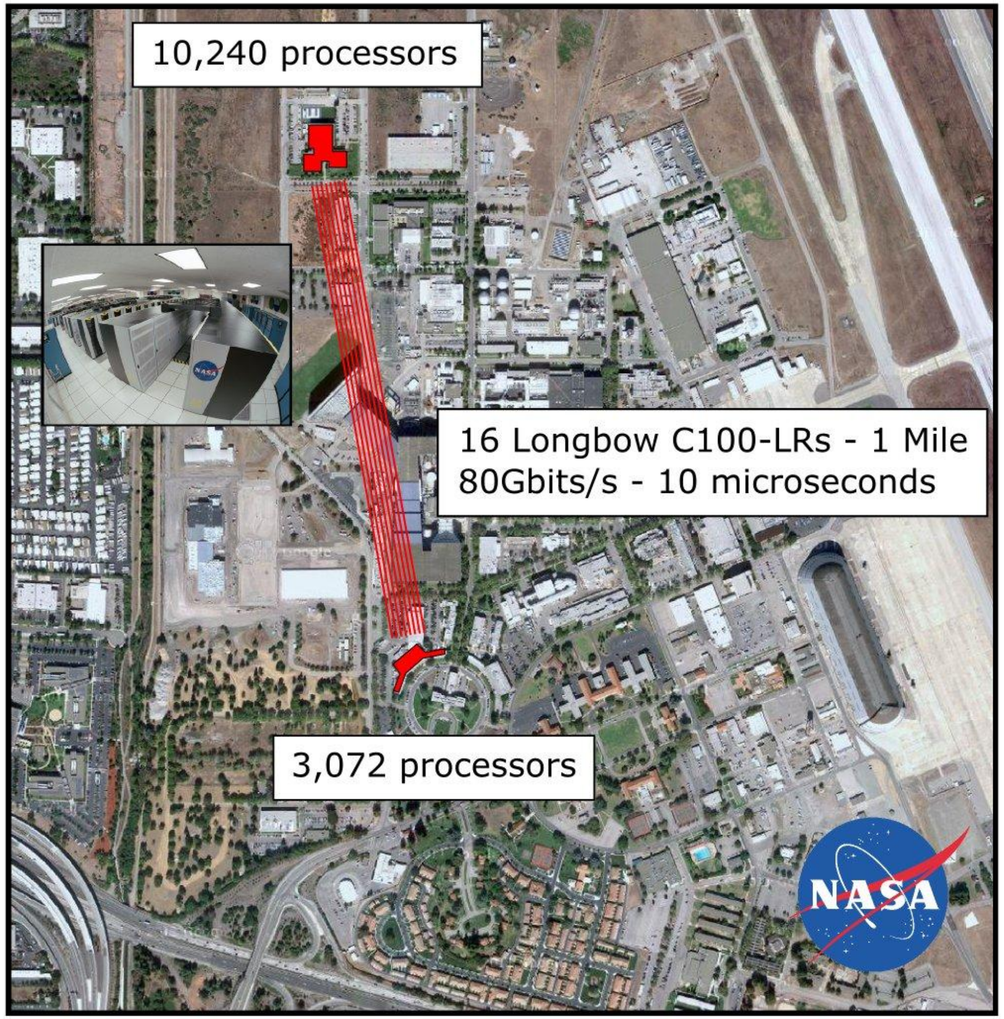

OBSIDIAN STRATEGICS

A*STAR Computational Resource Centre

InfiniCortex is like a living global brain

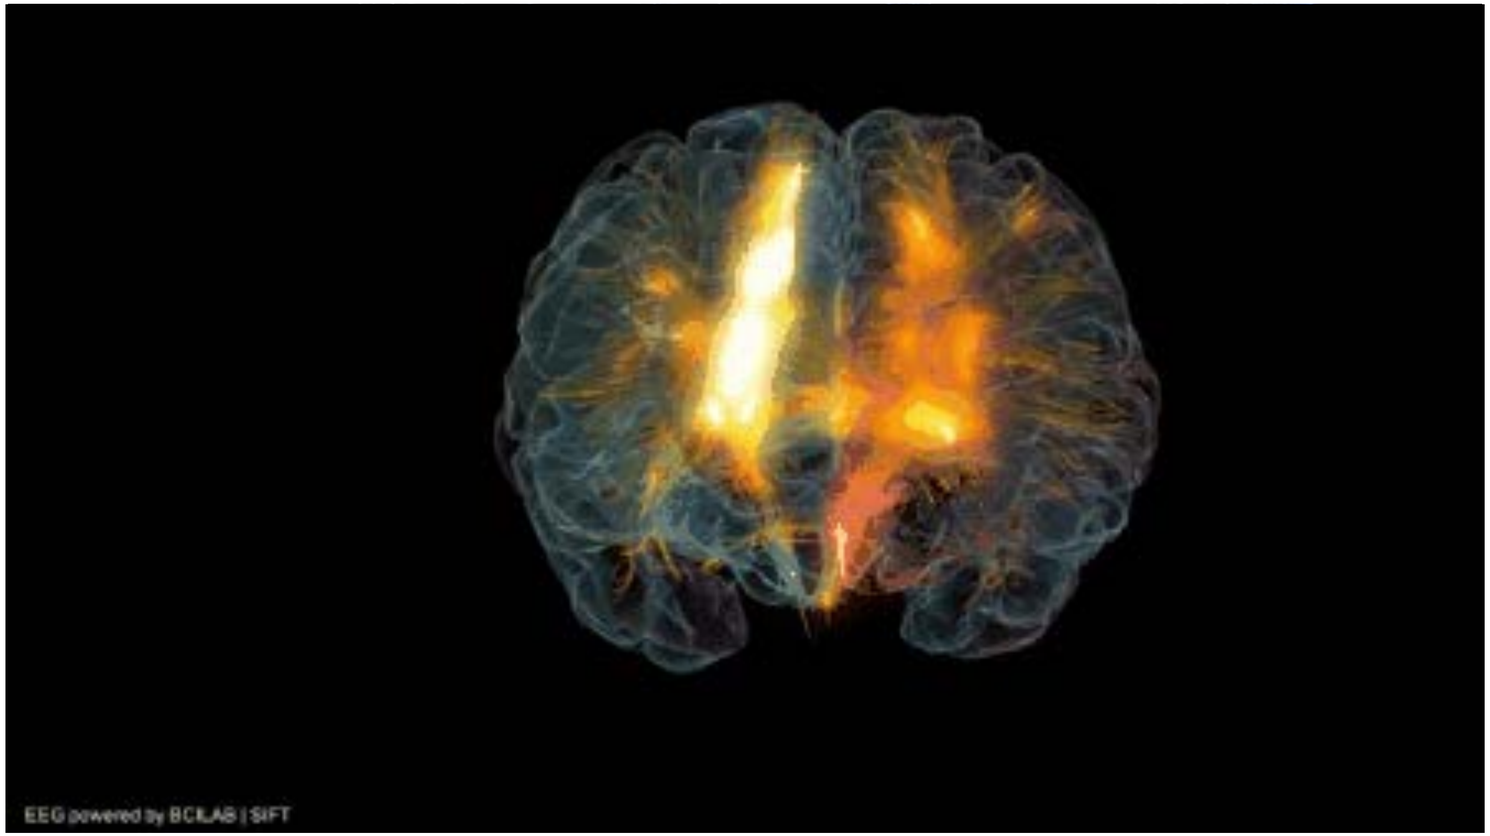

The **InfiniCortex** uses a metaphor of a human brain's outer layer, the **Cortex**, consisting of highly connected and dense network of neurons enabling thinking ….

to deliver **concurrent** supercomputing **across the globe** utilising trans-continental **InfiniBand** and **Galaxy of Supercomputers**

InfiniCortex is …

NOT GRID!

NOT CLOUD!

NOT “Internet”!

InfiniCortex Components

1. Galaxy of Supercomputers

- Supercomputer interconnect topology work
  by Y. Deng, M. Michalewicz and L. Orlowski
- Obsidian Strategics Crossbow InfiniBand router

2. ACA 100 & ACE 10

- Asia Connects America 100 Gbps, by November 2014
- Asia Connects Europe 10Gbps, established February 2015

3. InfiniBand over trans-continental distances

- Using Obsidian Strategics Longbow range extenders

4. Application layer

- from simplest file transfer: dsync+
- to complex workflows: ADIOS, multi-scale models

Galaxy of Supercomputers

- Supercomputers located at different geolocations connected into a ***Nodes of Super-Network (Super-Graph)***
- Supercomputers may have arbitrary interconnect topologies
- Galaxy of Supercomputers is a topological concept and is based on a topology with small diameter and lowest possible link number
- In terms of graph representation it may be realized as ***embedding*** of graphs representing Supercomputers' topologies into a graph representing the Galaxy topology

Investigators:
- Prof. Yuefan Deng, A*CRC & Stony Brook University
- Lukasz Orlowski, A*CRC & Stony Brook University
- Dr. Marek Michalewicz, A*CRC & Stony Brook University

32k5⊗32k5

Embedding of a 5-connected graph on 32 nodes into itself proves to be comparable to TOFU or 5D torus with equal or similar number of nodes.

Name of topology	*Number of nodes*	*Number of link*	*Diameter*	*Mean path length*
32k5⊗32k5	***1024***	***2640***	***9***	***6.31***
Tofu (6x5x3)	*1080*	*5400*	*9*	*5.04*
5D torus (4x4x4x4x4)	*1024*	*5120*	*10*	*5*
Tofu (4x4x8)	*1536*	*7680*	*11*	*5.67*

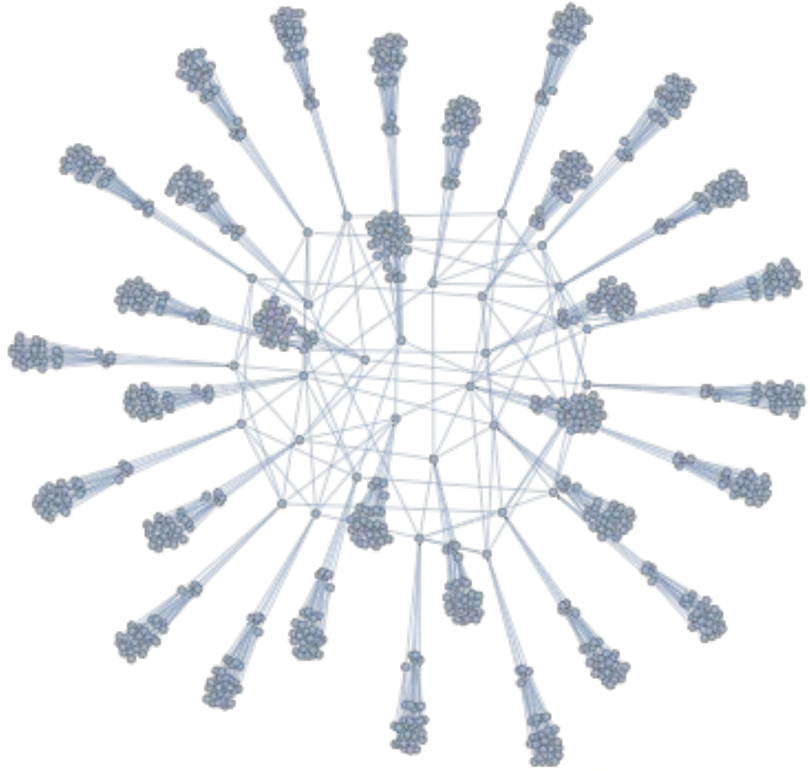

Galaxies of Supercomputers and their underlying interconnect topologies hierarchies

Lukasz P. Orlowski[1], Yuefan Deng[1,2,3] and Marek T. Michalewicz[1]

[1] A*STAR Computational Resource Centre, Singapore 138632, Singapore; [2] Stony Brook University, New York 11794-3600, USA; [3] National Supercomputer Centre in Jinan, Shandong Province, P. R. China

poster at ISC'14, Leipzig, June 2014

A*STAR Computational Resource Centre

Obsidian Strategics Hardware

Longbow C100 – 10G IB range extender
Dark Fiber/ xWDM Metro Area Networks (80 km)

Longbow X100 (Military) – 10G IB range extender/ router
10GbE/ OC-192, Wide Area Networks (unlimited distance)

Longbow E100 – 10G IB range extender/ router/ crypto
Dark Fiber/10GbE, Wide Area Networks (unlimited distance)

A-CWDM81 – 10G Optical Mux/Demux
Nine CWDM channels, Metro Area Networks (80 km)

Longbow C400 – 40G IB range extender
Dark Fiber/ xWDM Regional Area Networks (1~1,600 km)

Crossbow R400-6 – 40G
Six-port native IB router (enables multi-subnet LAN fabrics)

Computational Resource Centre

Ensuring Singapore is a Global Hub

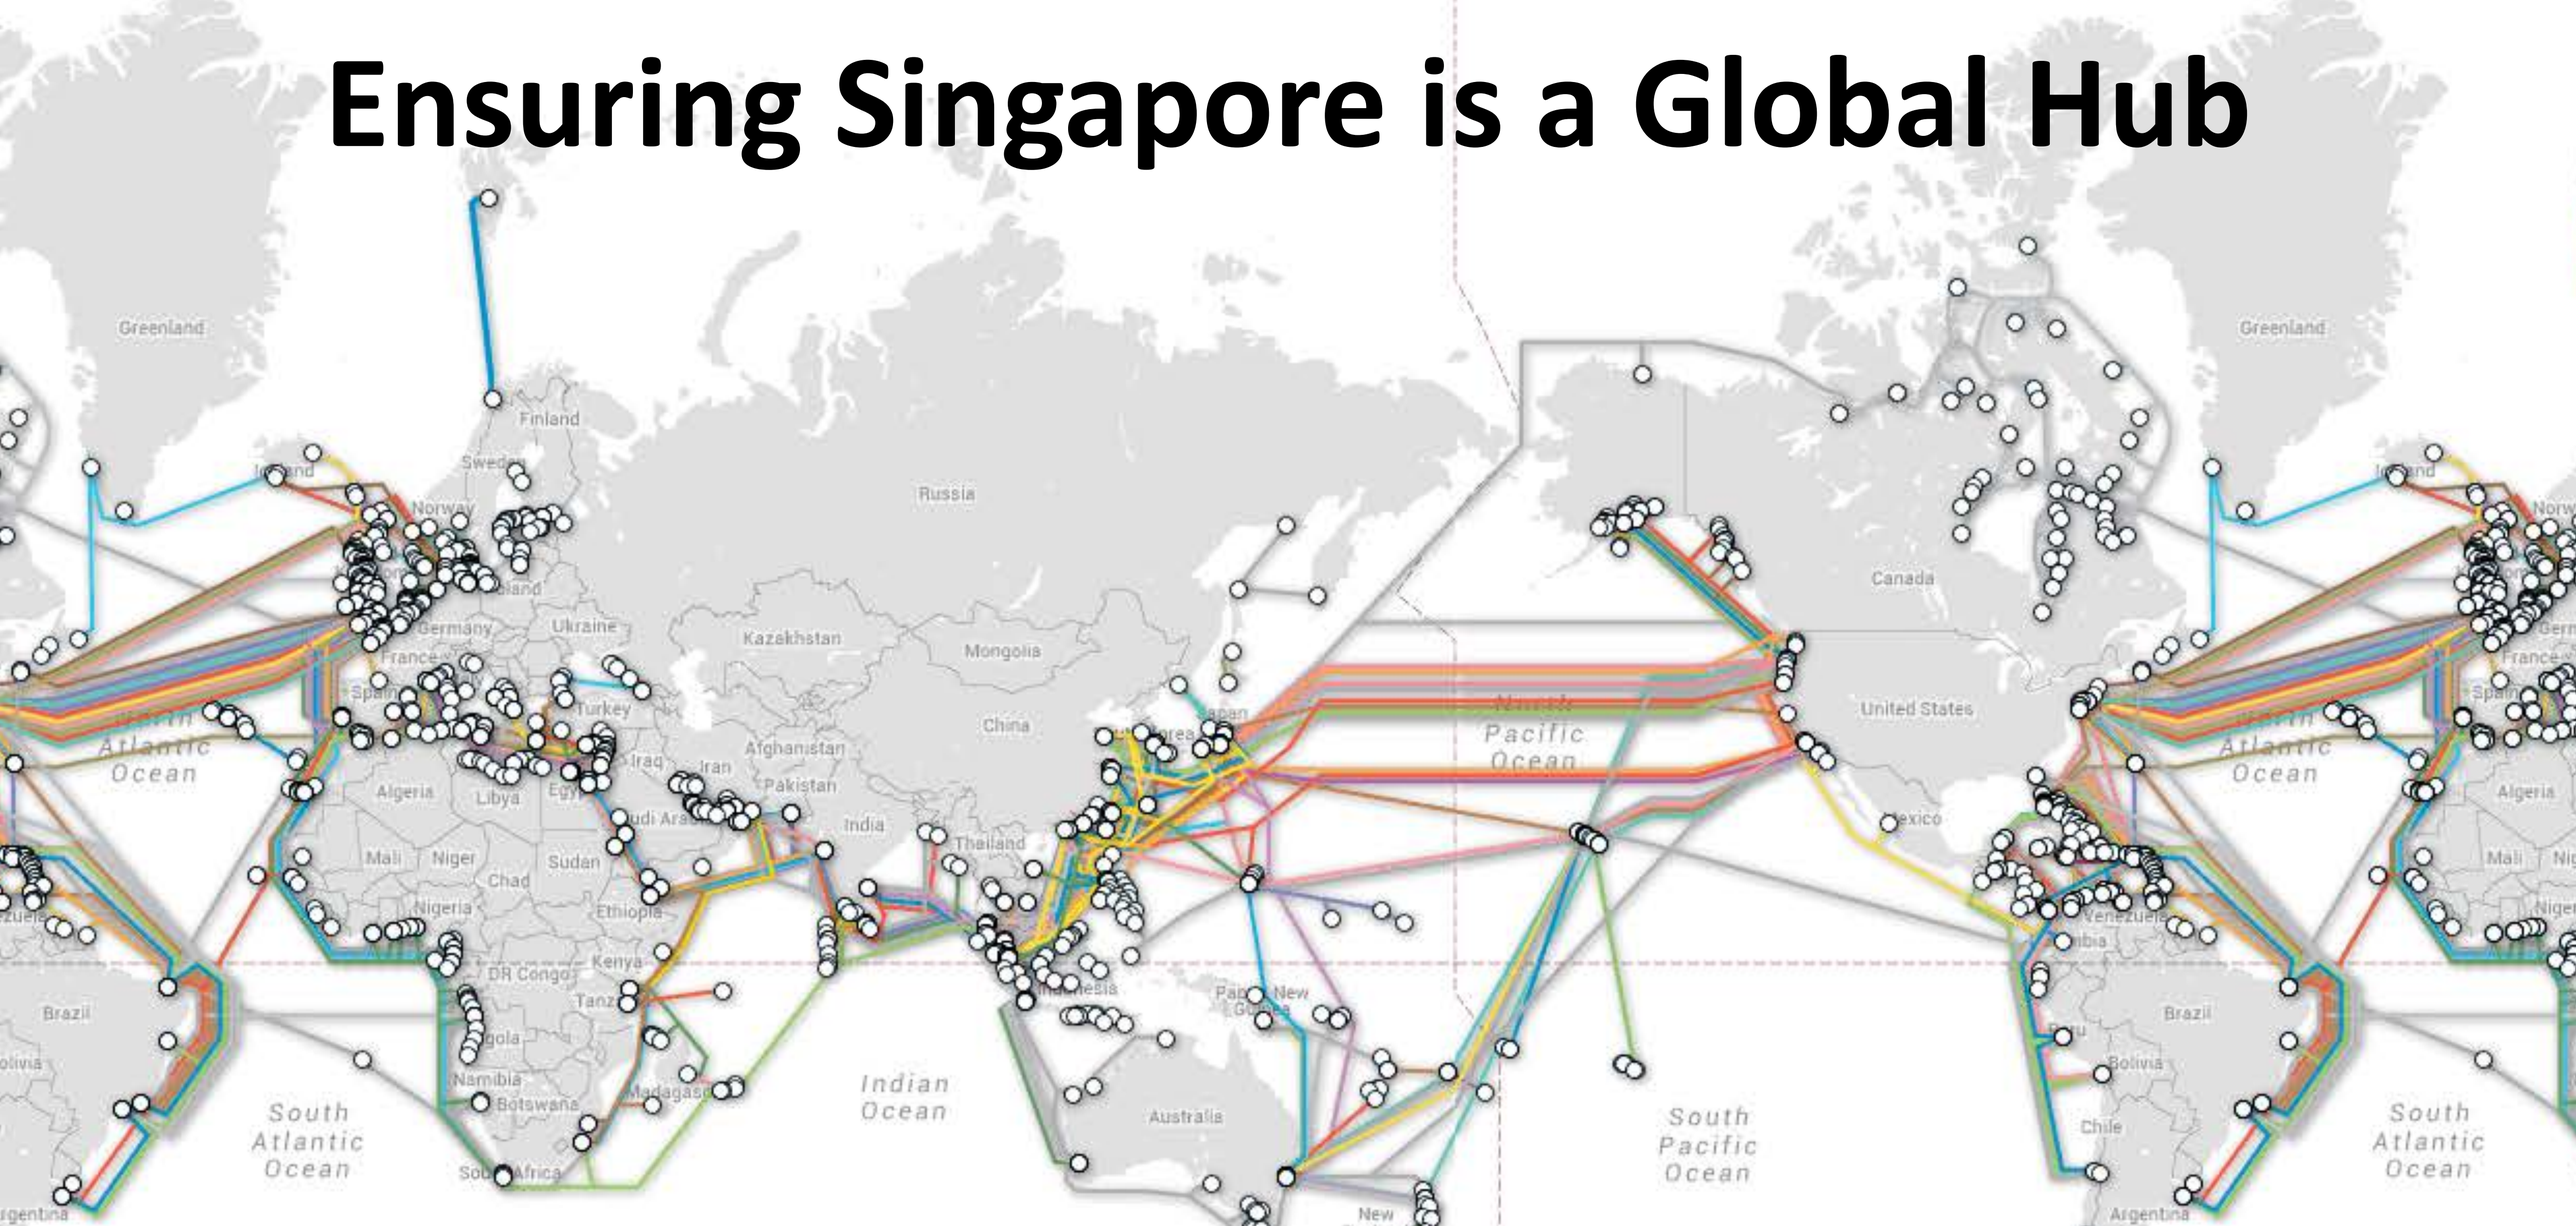

A National and International InfiniBand fabric linking Super Computers, Clouds and Data the foundation for a new level of global Science and Research

Extending InfiniBand beyond the Data Centre provides a more efficient way of using existing

InfiniCortex 2014 (phase 1)

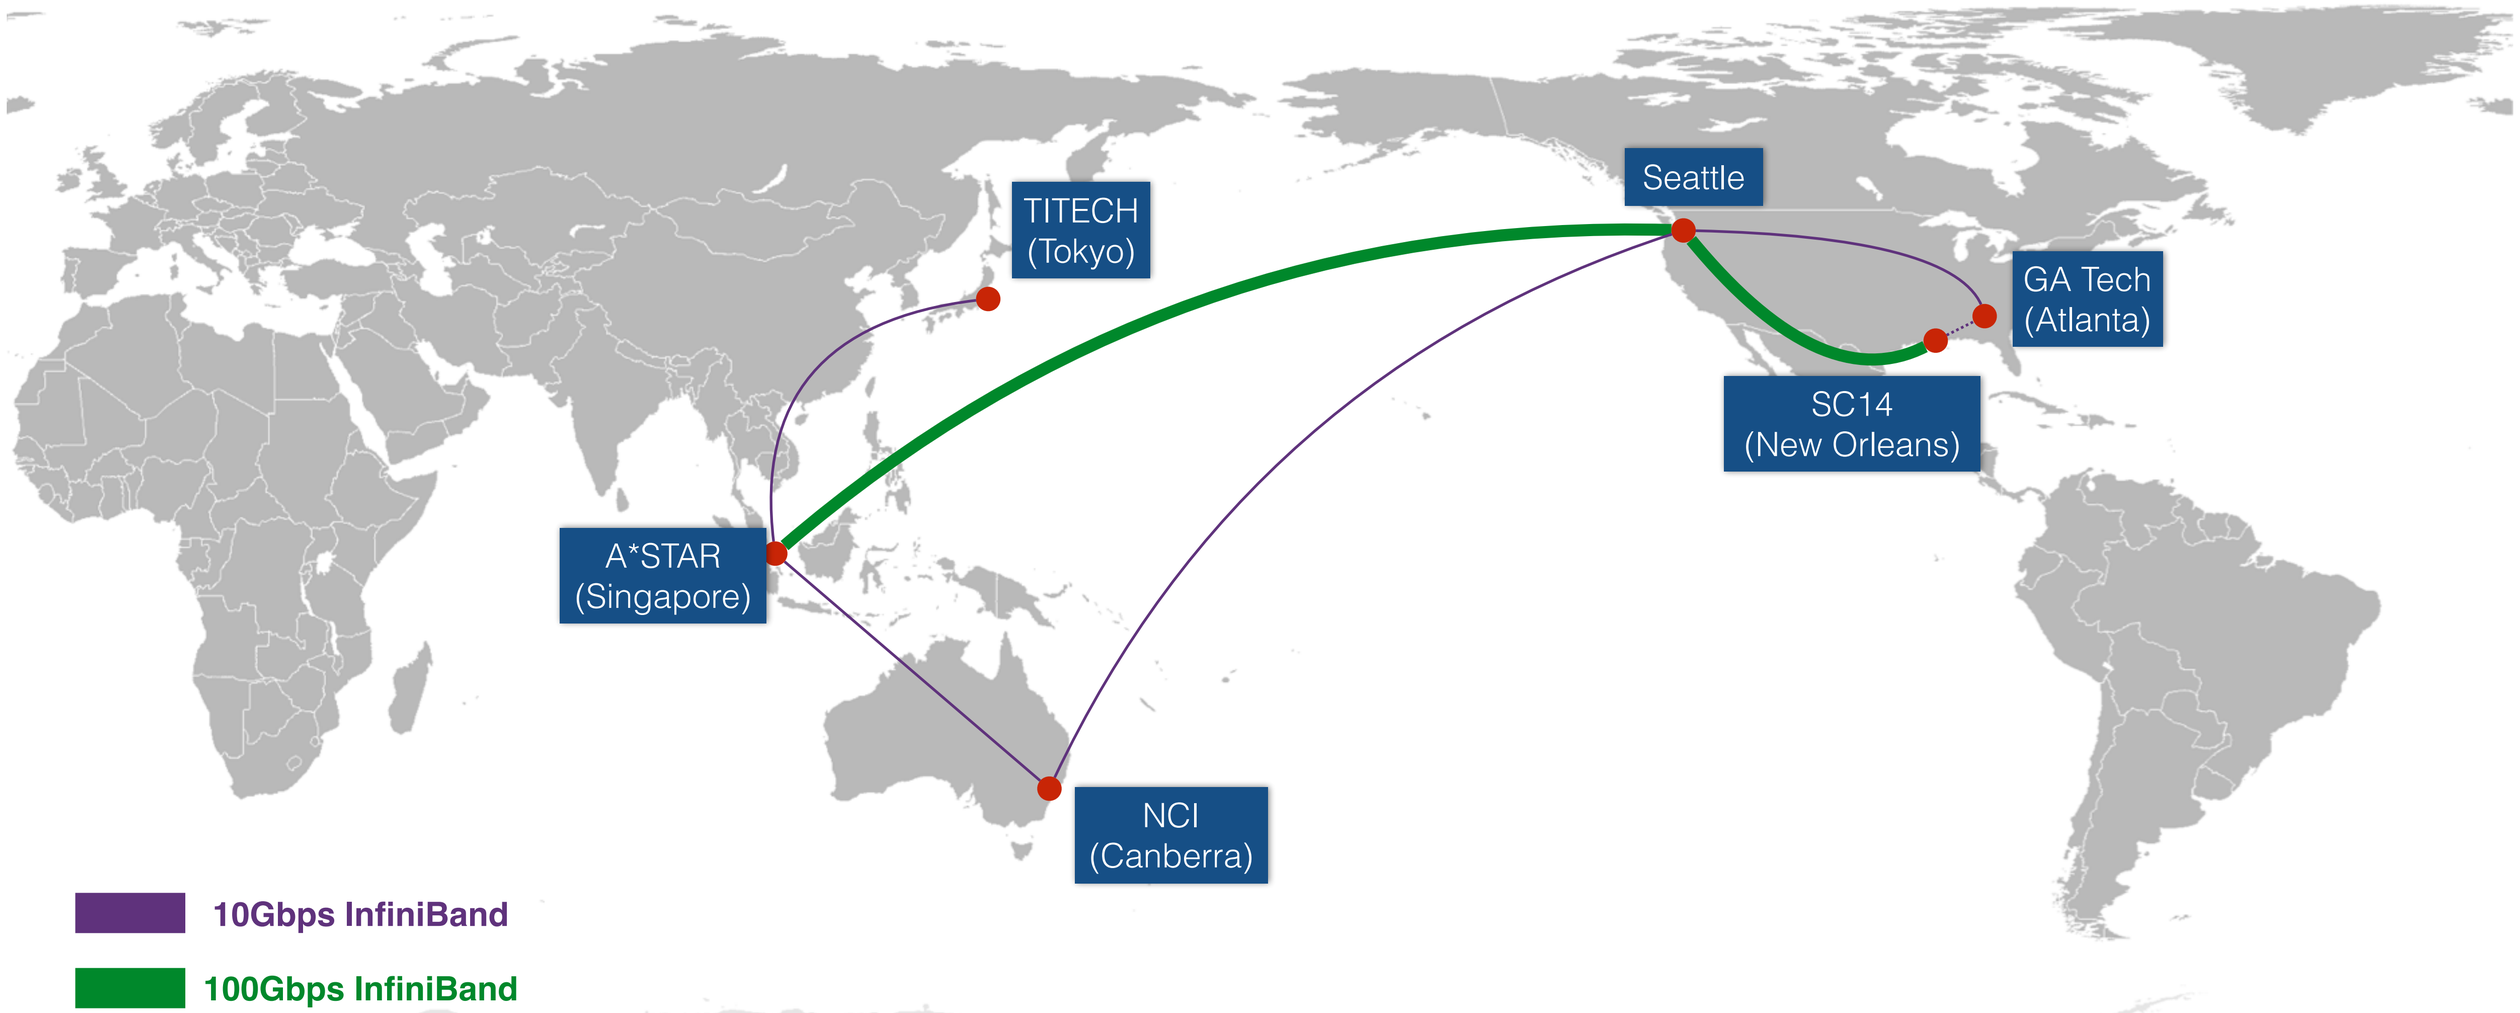

Enabling geographically dispersed HPC facilities to collaborate and function as ONE concurrent supercomputer, bringing the capability to address and solve grand challenges to the next level of efficiency and scale.

100Gbps Bandwidth Utilization

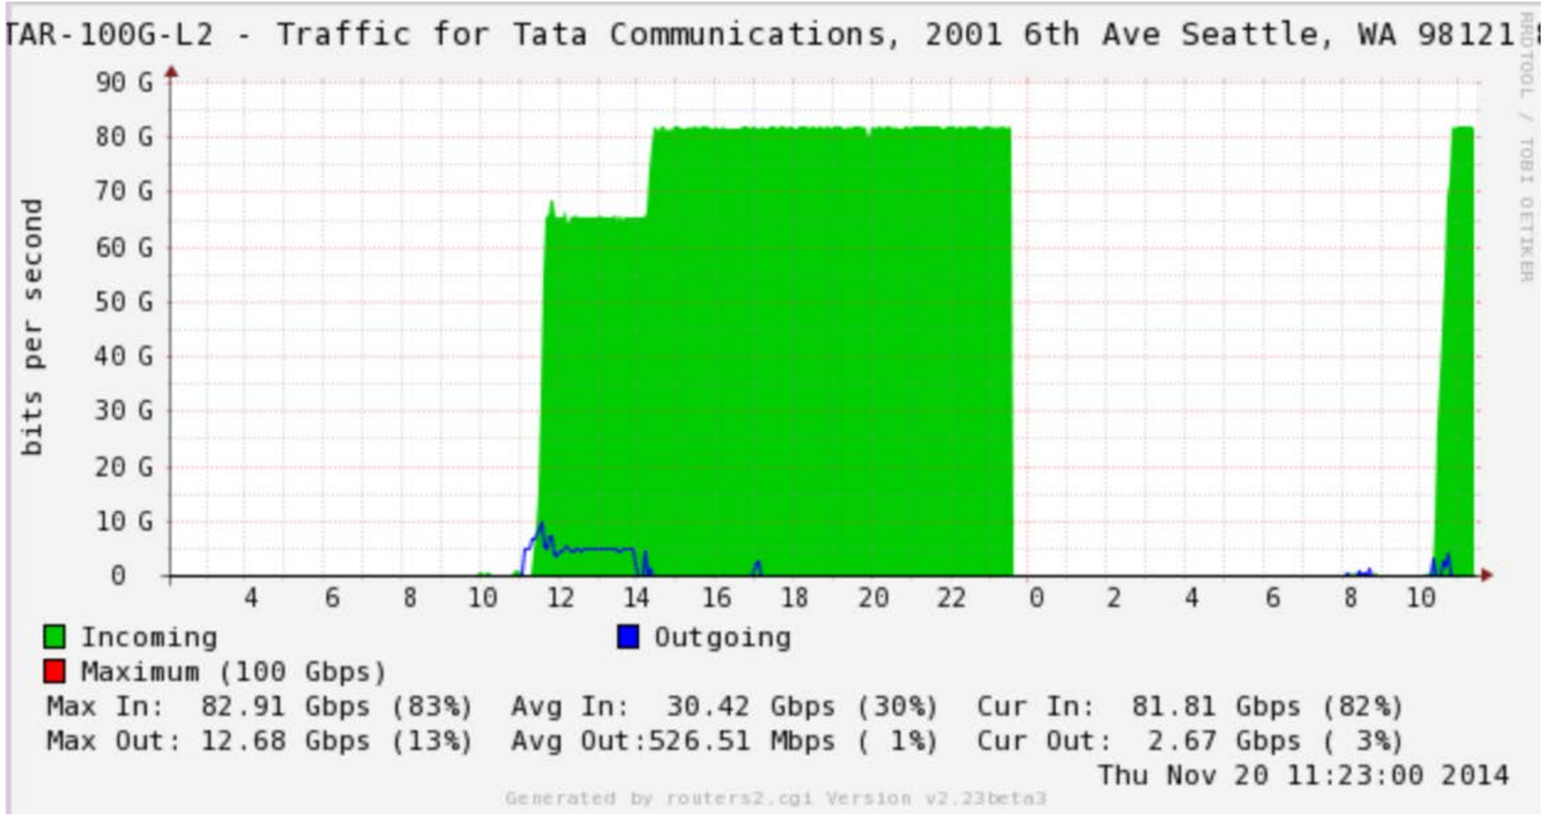

A*STAR Computational Resource Centre

Project Partners

Team - Singapore

Computational Resource Centre, A*STAR

SingAREN

Nanyang Technological University

Team - USA

Oak Ridge National Laboratory

Rutgers University

Georgia Tech

Stony Brook University

ESnet

Internet2

SCinet

Team - Japan

東京工業大学 Tokyo Institute of Technology

NII National Institute of Informatics

NICT

TEIN Cooperation Center

Team - Australia

NCI

aarnet

Commercial Partners

Obsidian Strategics

TATA COMMUNICATIONS

DataDirect Networks

ciena

intel

Mellanox Technologies

sgi

Commercial Sponsors

BlueTel

Pacific Wave

CenturyLink Government

3D Networks

cisco

Juniper Networks

InfiniCortex 2015 (phase 2)

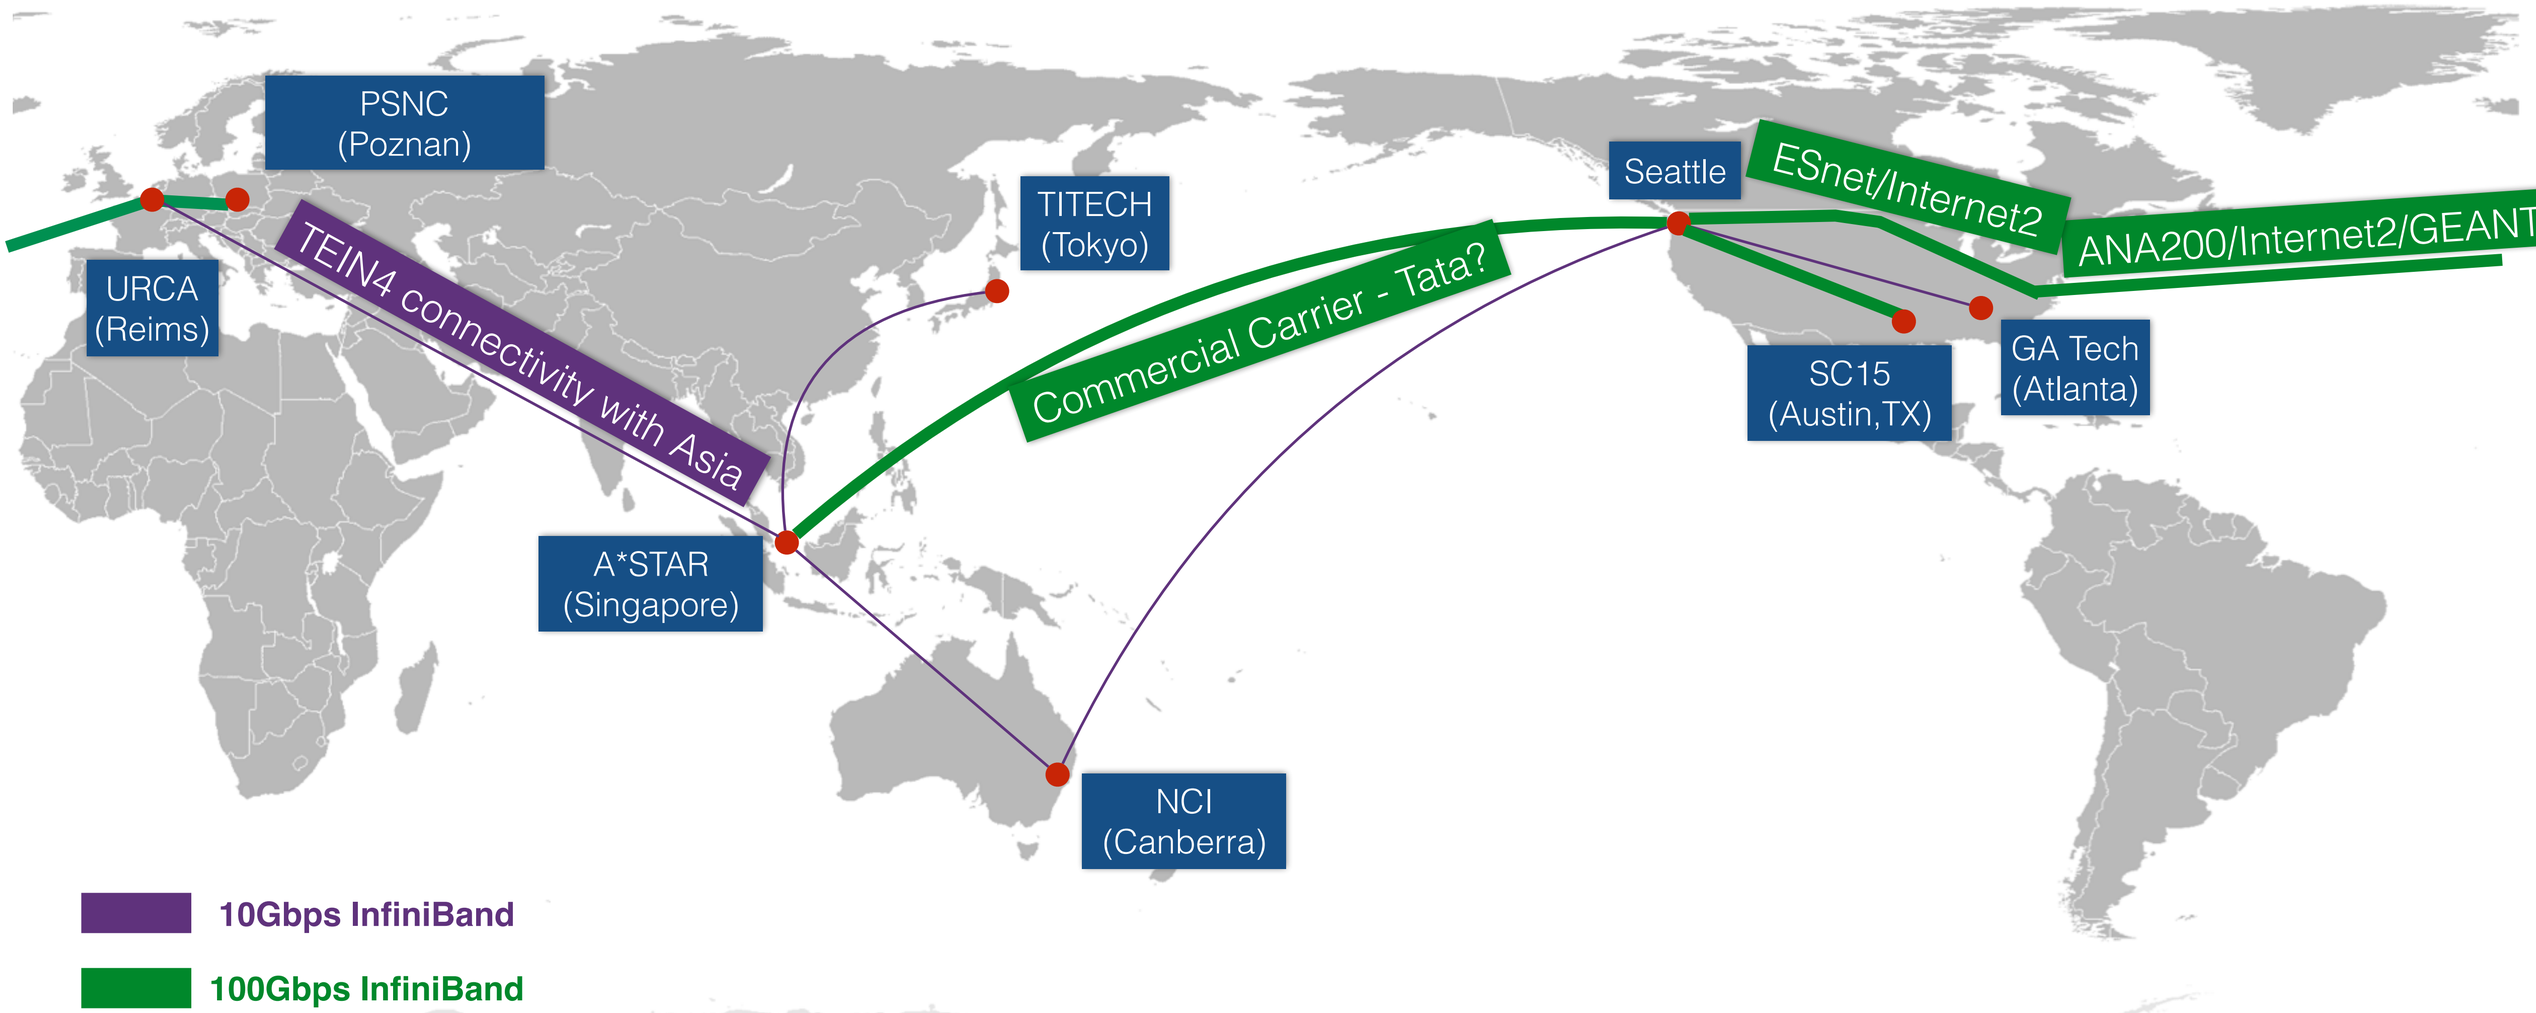

100Gbps InfiniBand East-ward link: Singapore-trans-Pacific-USA-trans-Atlantic-Europe

10Gbps InfiniBand West-ward link: Singapore-Europe (via TEIN4)

InfiniCortex 2015 (Europe)

2015 European Partners:

University of Reims, France
University of Lille, France
Poznan Supercomputing and Networking Centre, Poland
ICM Warsaw, Poland

Circuit providers:
GEANT, RENATER, PIONEER

Singapore InfiniBand connectivity

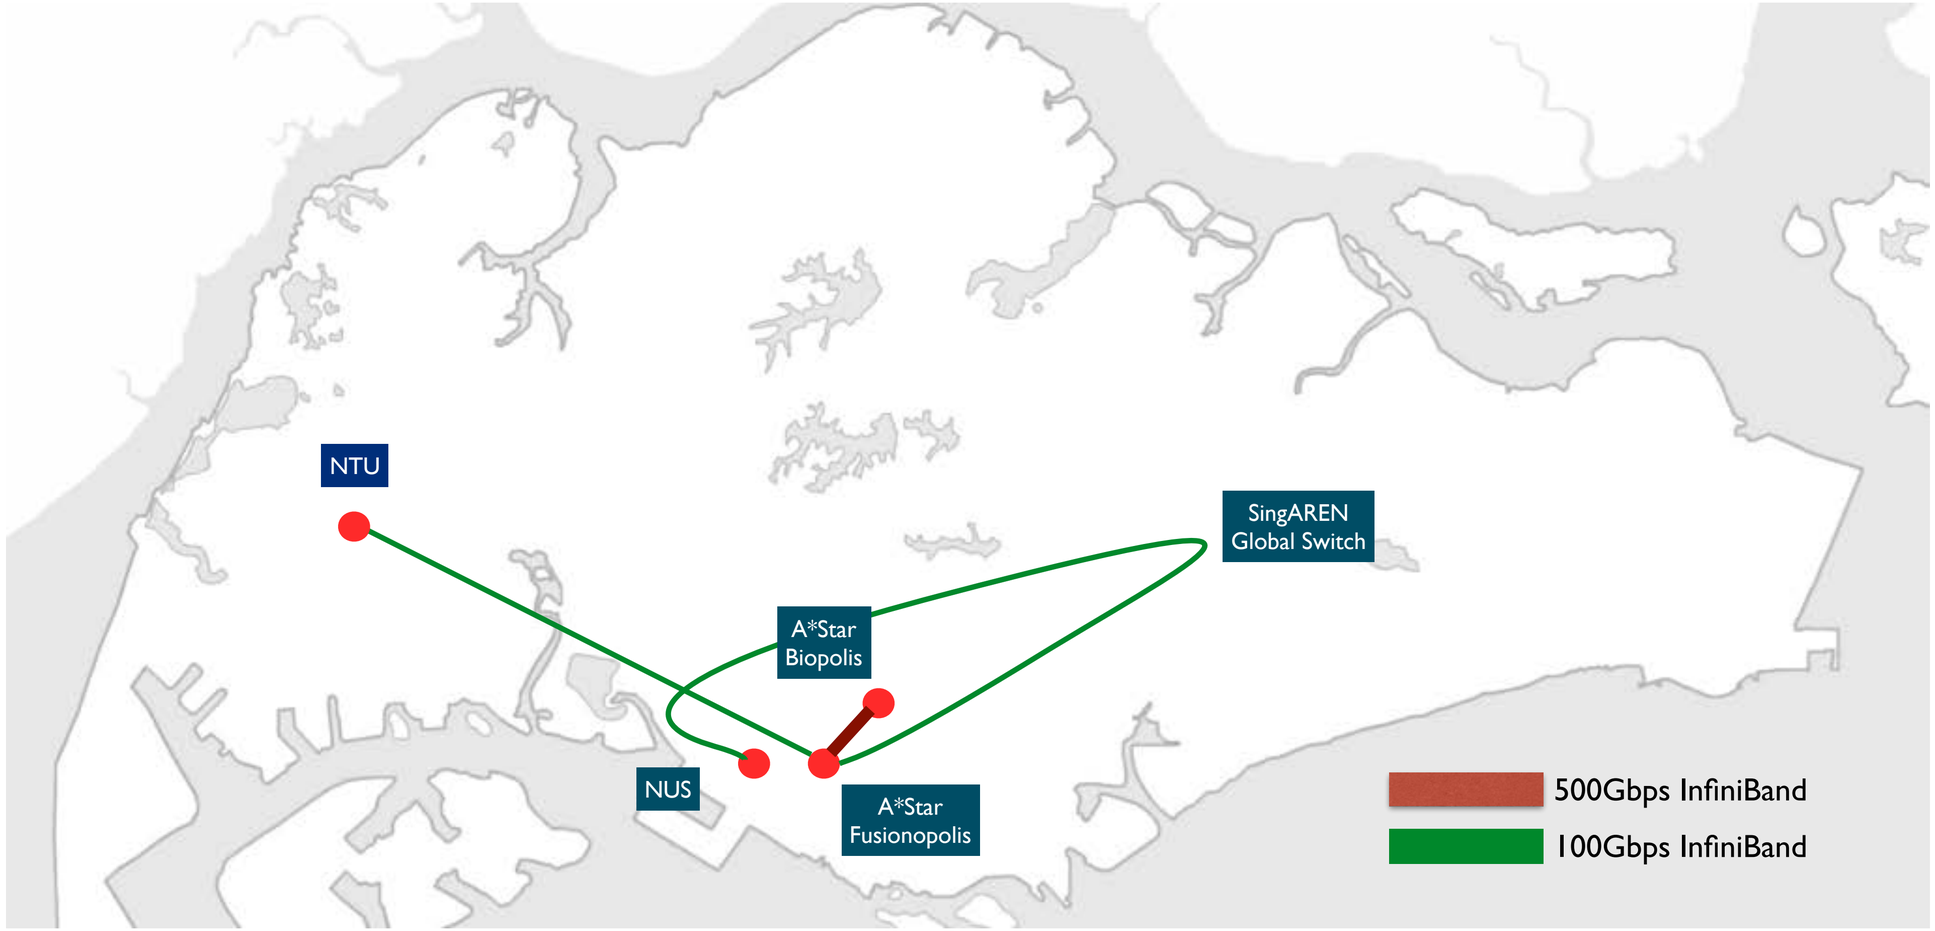

Connecting all National Supercomputing Centre stakeholders: A*STAR, NUS, NTU and others with 100Gbps + InfiniBand links.

A*STAR Computational Resource Centre

Singapore InfiniBand connectivity

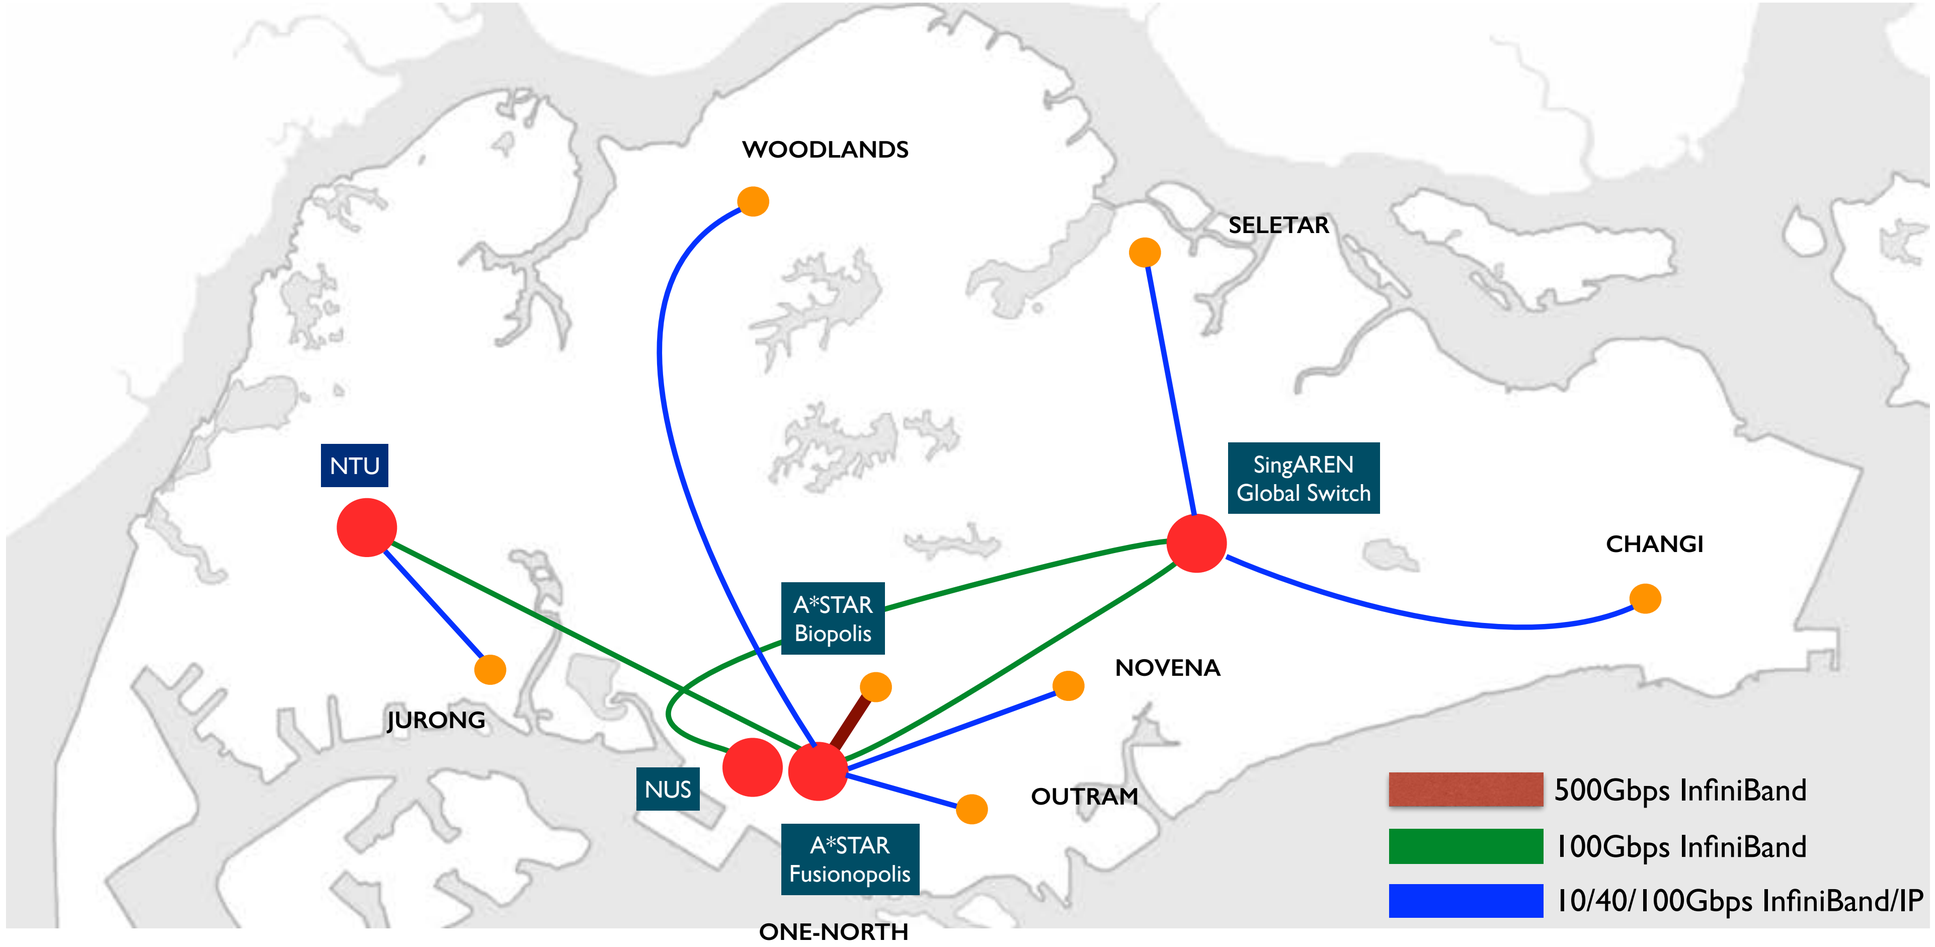

Connecting all National Supercomputing Centre stakeholders: A*STAR, NUS, NTU and others with 100Gbps + InfiniBand links.

A*STAR Computational Resource Centre

Streaming Segmentation of Large Pathology Tissue Images

Abstract:
Demo showed nuclear segmentation on a high resolution whole slide tissue color image (53,000x36,000 pixels, RGB) stored on a cluster at ACRC in Singapore.

The image is partitioned into tiles. The tiles that contain enough tissue data are streamed from multiple cluster nodes via the ORNL ADIOS system over Longbow connections to a cluster at Georgia Tech (GT) in the US. The tiles are processed on the GT cluster using the SBU RT middleware as they are received to segment nuclei. Segmented nuclei and the tiles are assembled into an image.

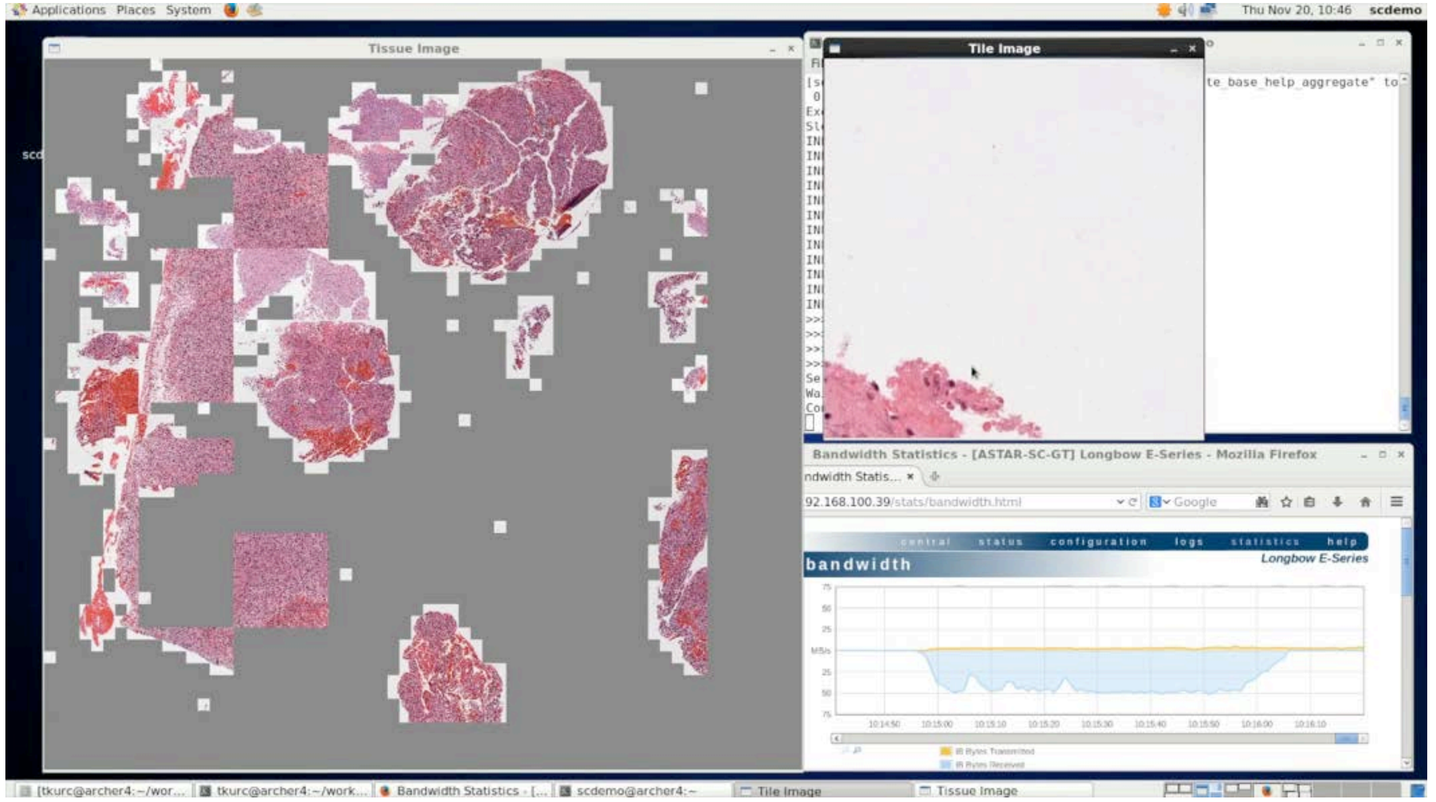

Team:

Tahsin Kurc, SBU and ORNL Scott Klasky, ORNL Jong Choi, ORNL Joel Saltz, SBU

Stony Brook University SDG ADIOS OAK RIDGE NATIONAL LABORATORY
MANAGED BY UT-BATTELLE FOR THE U.S. DEPARTMENT OF ENERGY

Near Real Time (NRT) Application Scenario: Plasma Disruption Detection

- Plasma disruption
  - Lead to the loss of stability and/or confinement of tokamak plasmas
  - Cause fast thermal and/or current quench
  - Could damage multi-billion tokamak

- The experimental facility may not have enough computing power for the necessary data processing

- Distributed in transient processing
  - Make more processing power available
  - Allow more scientists to participate in the data analysis operations and monitor the experiment remotely
  - Enable scientists to share knowledge and process

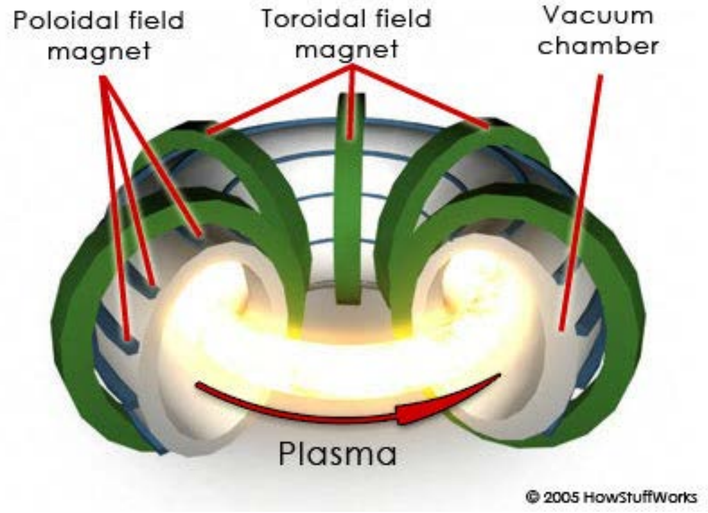

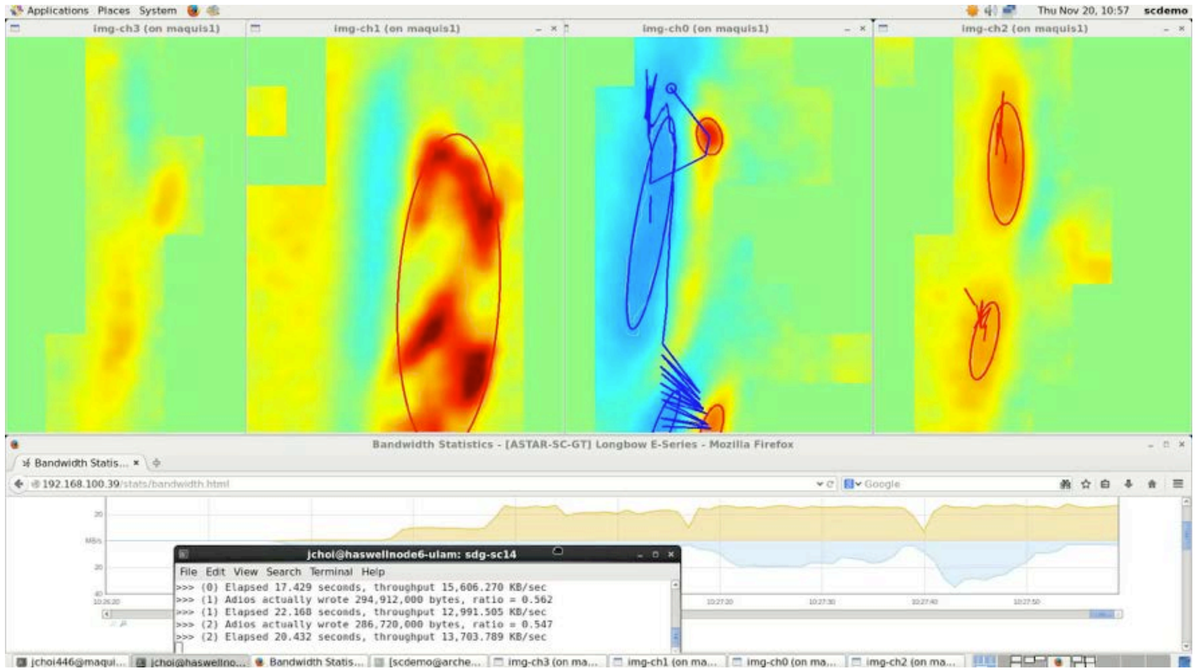

Genomics and Diagnostics

HPC clouds
Reference Mapping

Australia
NCI
NATIONAL COMPUTATIONAL INFRASTRUCTURE

Emerging Supercomputing Technologies

1. INTERCONTINENTAL offsite supercomputing
2. Integrating HPC clusters and HPC clouds
3. Making distant IB networks to be local ones
4. Super fast data transportation (0.83Gb/s)
5. Solutions to migrate and analyse BIG data
6. Proof-of-concept by a Bioinformatics NGS workflow

NGS reads

SAM
(Sequence alignment/map)

Singapore
Computational Resource Centre
A*STAR

SAM files

BAM
(Binary alignment/map)

SC14
New Orleans, LA | hpc matters.

NGS
(Next-Generation Sequencing)

Analysis
Visualization

The InfiniCortex Application Layer

Genomics and Diagnostics

Up to 1.2TBytes genomics files sent from Australia to Singapore in less than 30min for processing on very large shared memory machine

Work done with National Computing Infrastructure, ANU, Canberra, Australia.

Distance ~30,000km

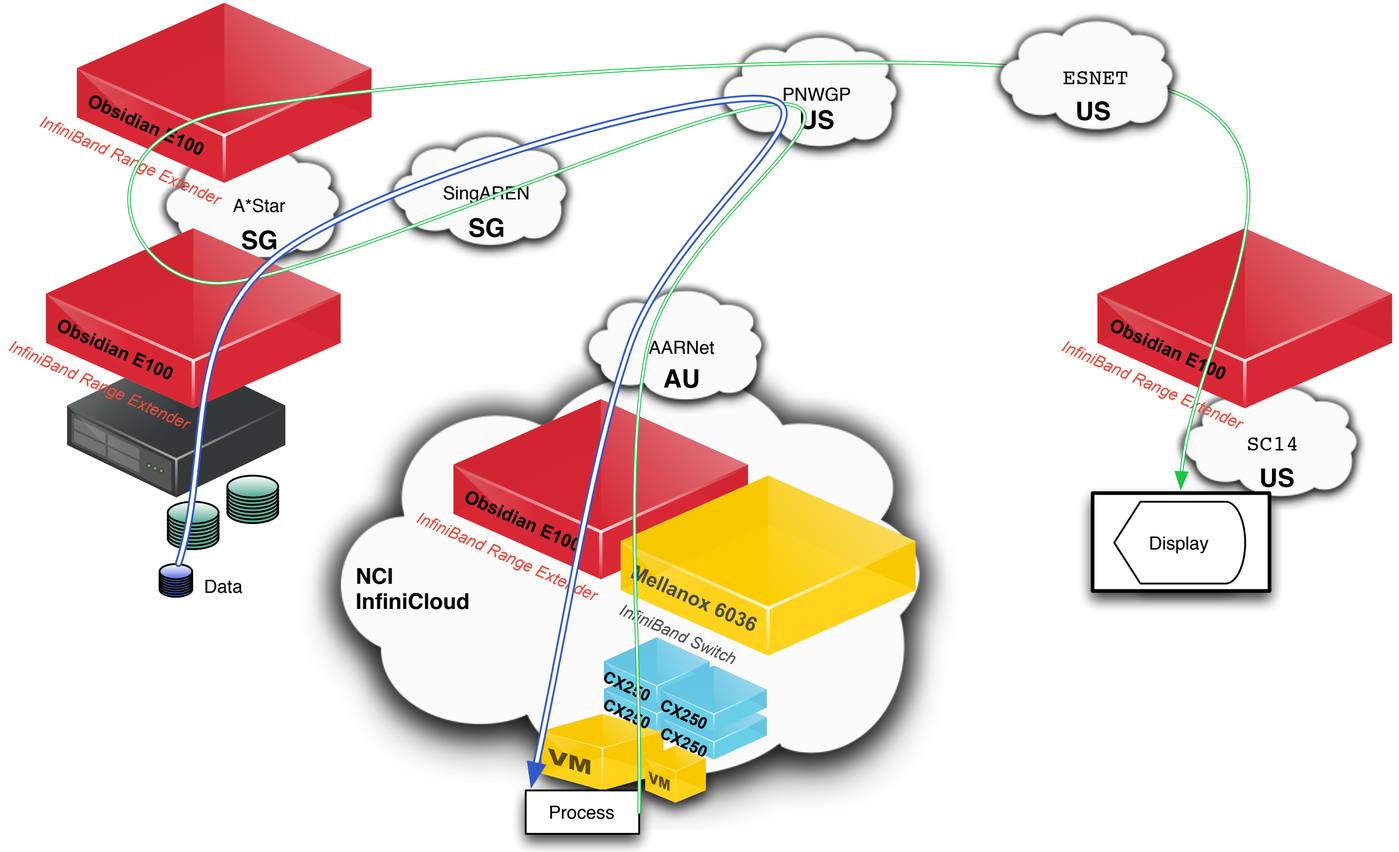

Computational Resource Centre
A*STAR

SG to AU data transfer speed

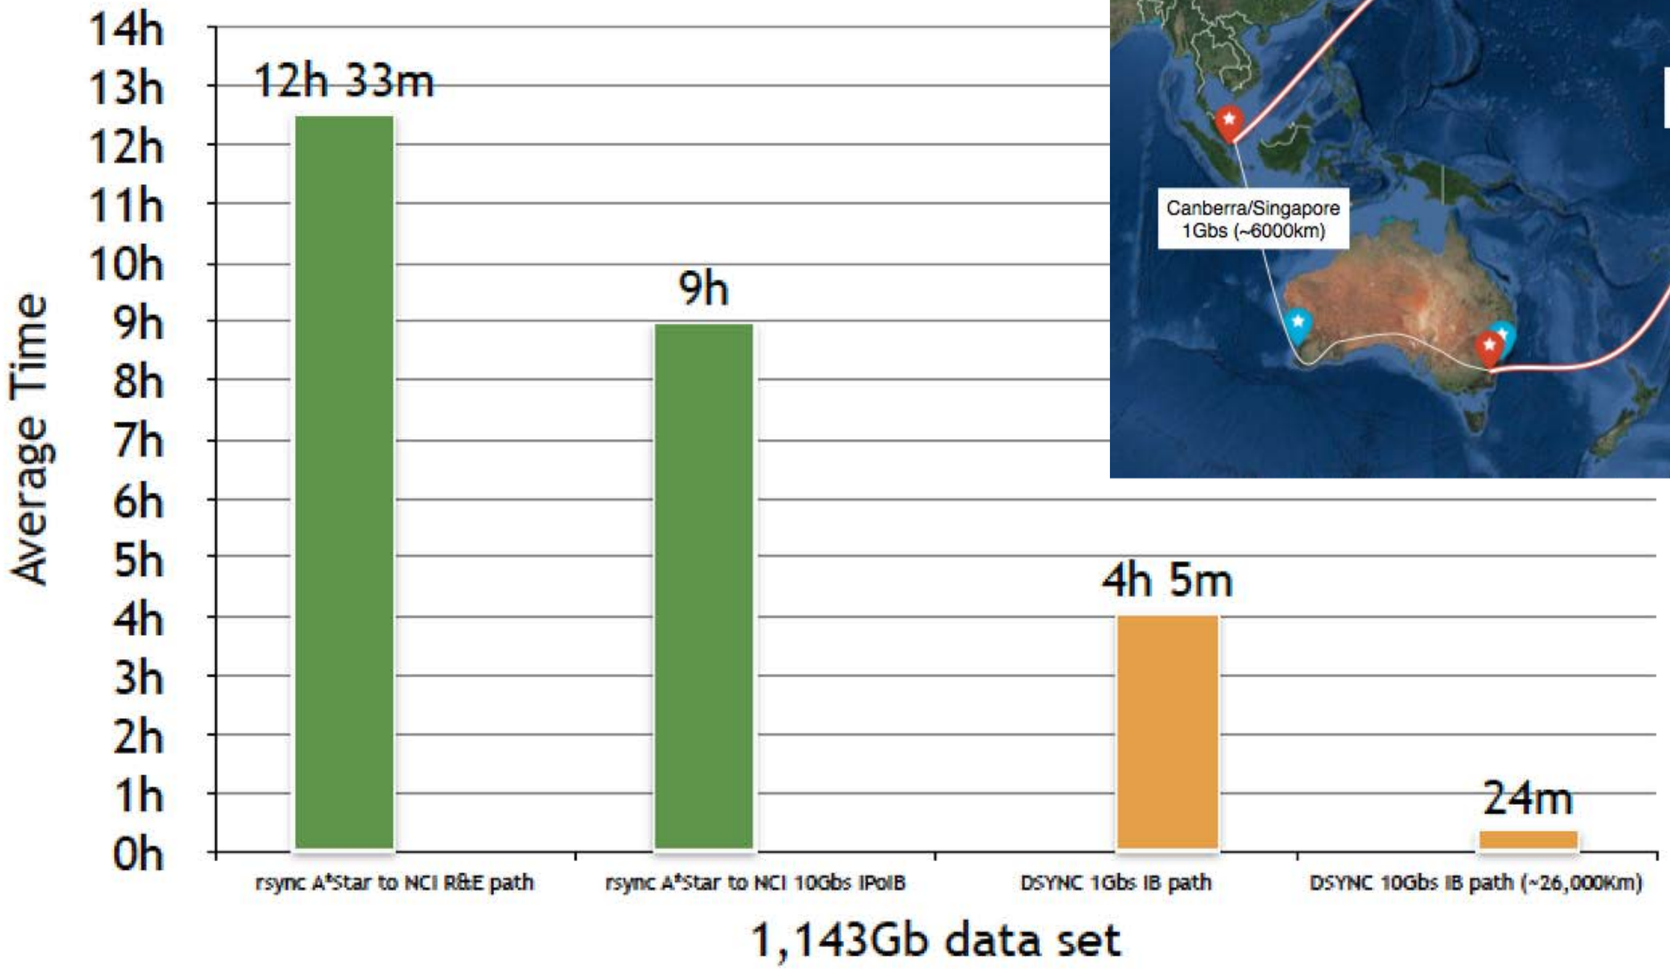

Awards

- Ministry of Trade and Industry (MTI) 2015 Gold Award for Innovative Project
  - 2015 A*STAR Innovation Award

  - FutureGov Singapore Award 2015 in Technology Leadership category
    - CIO 100 HONOUREE 2015

A*STAR Computational Resource Centre

Initiatives

- test Infiniband routing using BGFC
- test InfiniCloud (provision VMs across continents w/ Infiniband support)
- test global on demand HPC using Bright Cluster Manager
- test genomic databases replication using Garuda framework
- test extremely fast I/O (DDN Infinite Memory Engine)
- test DDN Web Object Scaler (WOS) Storage for bioinformatics and genomics data storage
- test applications in the field of molecular docking and linear systems solving

BGFC: A New Subnet Manager

- a 100% re-imagined subnet manager to address limitations of OpenSM
- no code borrowed from OpenSM
- mathematically perfect routing algorithm
- support for native InfiniBand
- handles compound topologies
- persistent database to avoid large scale re-computations
- high quality and deterministic connectivity metrics
- performance-sensitive runtime core written in C++11 (leveraging BGL-*BOOST Graph Library*)
- correctness-sensitive support code written in python-rdma (single language re-implementation of the entire OFED diagnostics tool set
- allows Infiniband routing between different subnets in conjunction with Obsidian R400 InfiniBand router
- supports different topologies in every subnet

InfiniCloud: NextGen Platform for Data Intensive Science

- true HPC cloud
- high throughput global RDMA (up to 100Gpbs) for borderless data intensive computing
- well suited for bioinformatics and computational genomics

InfiniCloud

DNA → RNA → Protein → Metabolites

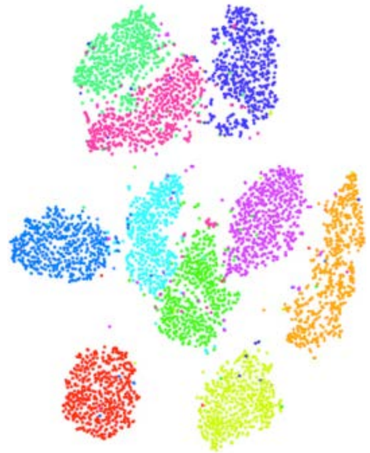

Function
Mechanism

Diagnosis
Treatment
Prevention

Comparing 'omics profiles among different states → Finding patterns Finding predictors → Insights for basic and translational science

A*STAR Computational Resource Centre
A*STAR Institute of Molecular and Cell Biology
NCI National Computational Infrastructure
NUH National University Hospital A member of the NUHS
NUS National University of Singapore Yong Loo Lin School of Medicine

Global on Demand HPC

Importance of a cluster manager for InfiniCortex project

- commission and decommission nodes remotely (via PXE boot)
- deploy, update and synchronise system image of nodes
- unify user account policies and credentials, software stacks, job scheduling policy

Experiment

- head node in Singapore remotely provisions a server in NCI via GE and Infiniband

Progress

- a remote node in NCI was PXE booted by the head node in ACRC
- OS image was deployed through a 1GB private ethernet link provided by Longbow management port
- Bright Cluster GUI provisions the NCI node from A*CRC

Future Work

- remote boot via IB instead of private GE
- provisioning cloud incidents on remote nodes

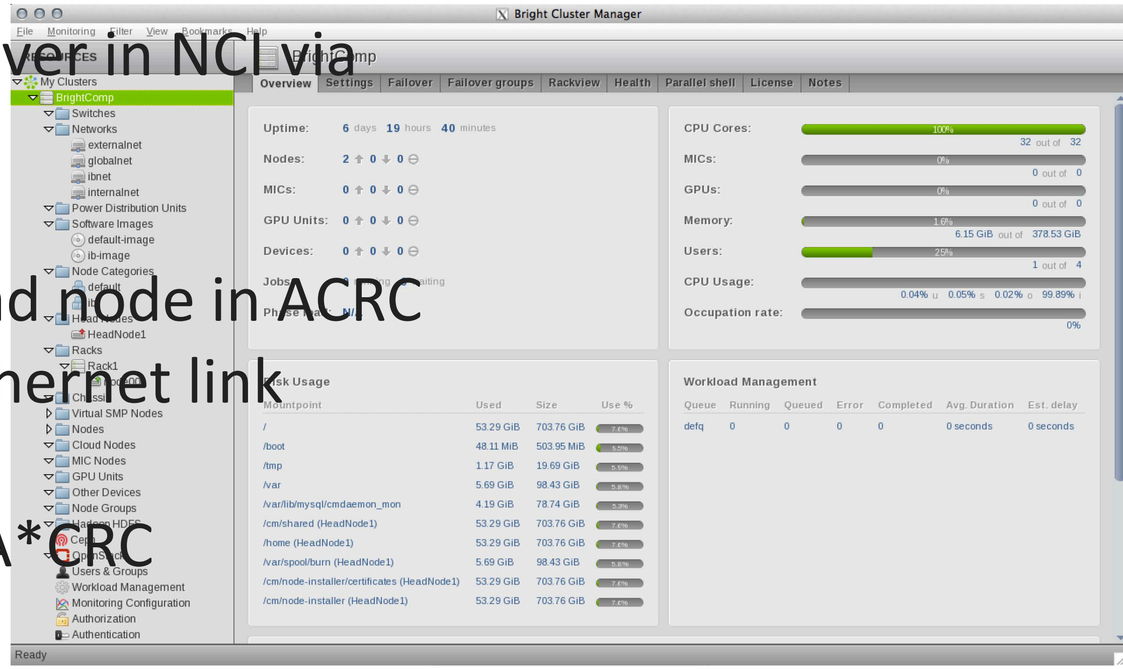

Computational Resource Centre A*STAR

NCI NATIONAL COMPUTATIONAL INFRASTRUCTURE

Bright Computing

Genomic databases replication - Garuda

Garuda – The way biology connects

- open, community-driven, common platform that provides a framework to interface, discover, & navigate through different applications, databases and services in bio-medical research
- provides language independent API to connect software as gadgets (over
- explore them through the gateway and operate them through the dashboard
- supported by a global alliance of leaders in computational biology and informatics

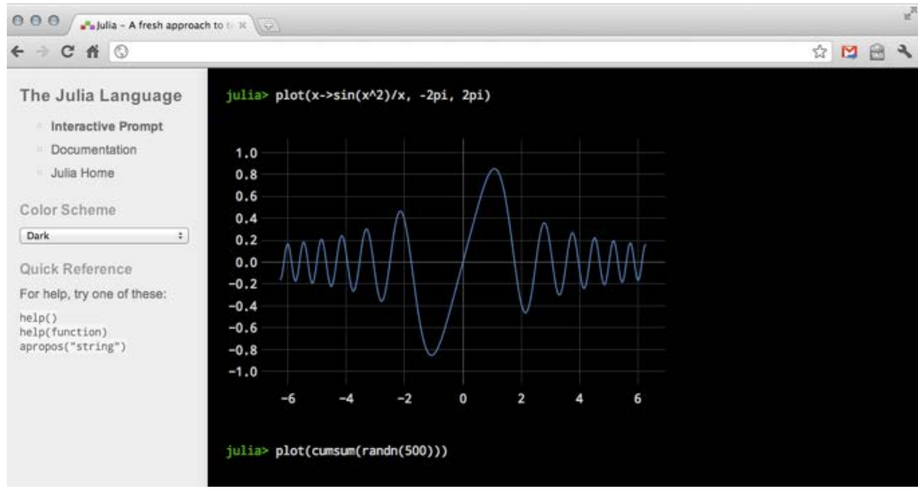

Project scope:

- run Julia (high-level, high-performance dynamic programming language for technical computing) in tight integration with Garuda

Computational Resource Centre A*STAR
SBI The Systems Biology Institute
julia
GARUDA THE WAY BIOLOGY CONNECTS

Infinite Memory Engine

- IME unleashes a new I/O provisioning paradigm
- software defined storage application - transparent, extendable, non-volatile memory (NVM), that provides latency reduction and greater bandwidth and IOPS performance for scientific, analytic and big data applications
- general availability mid-2015 for sites with a clustered compute environment and run a parallel file system, such as Lustre or GPFS

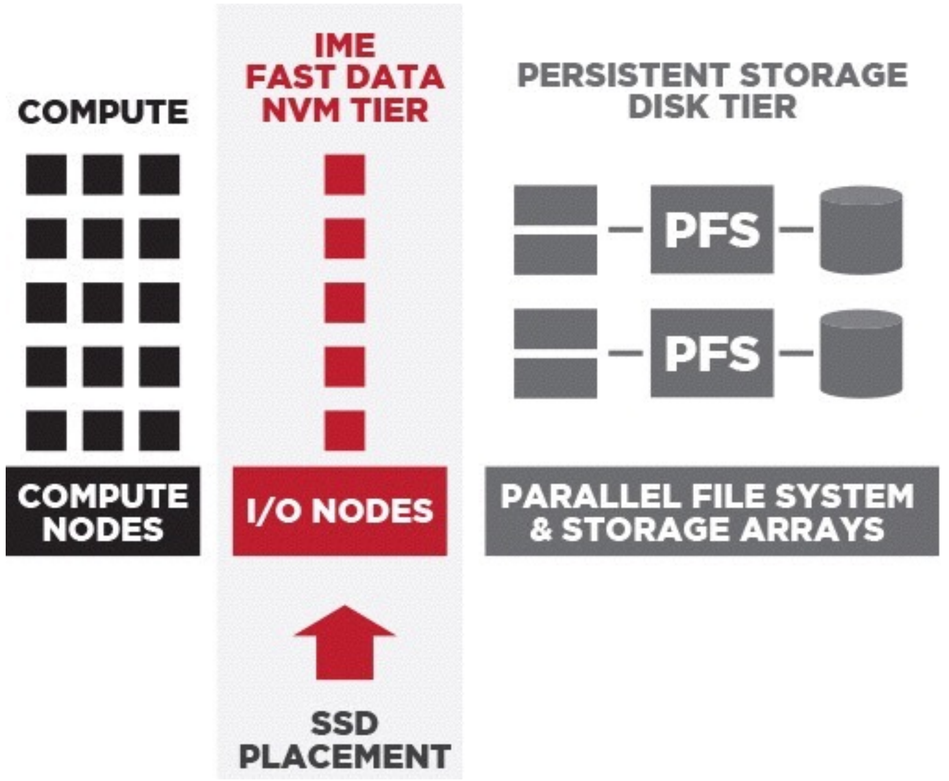

Bioinformatics and Genomics Data Storage

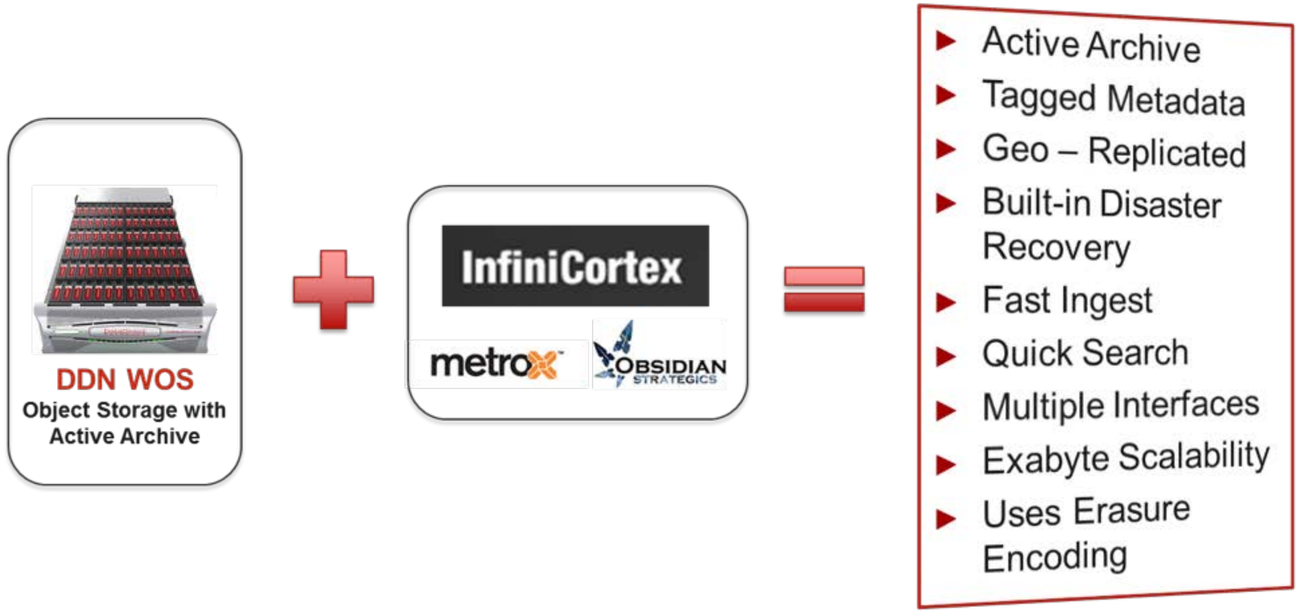

- DDN Web Object Scaler (WOS) Storage - high-performance object storage platform designed to easily store PB of unstructured data with the highest availability
- WOS scales to hundreds of PB (in clusters as large as 30PB) across multiple, geographically distributed WOS sites to present a globally accessible, single federated namespace
- Object Storage:
  - backend architecture designed as a single pool of storage nodes (no file system hierarchy)
  - when objects are stored, an identifier is created to locate the object in the pool
  - apps will use identifier to retrieve the right data through a REST API

A*STAR Computational Resource Centre

DataDirect NETWORKS Information in Motion

Maison de la Simulation Reims

Inverse docking method for new proteins targets identification

- scale docking method to a multi-cluster infrasctructure using a Reims web portal dedicated to life-science simulations

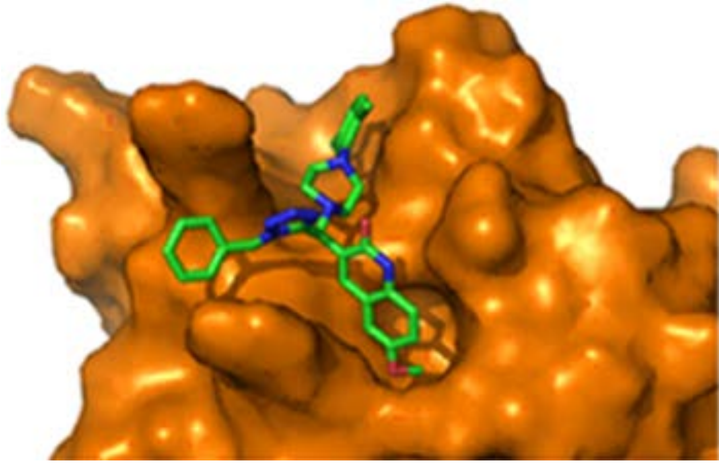

Realtime visualization of valorisation of agricultural resources

- CPMD simulation of a parallelized plane wave/pseudopotential implementation of Density Functional Theory, particularly designed for ab-initio molecular dynamics

An energy efficiency study on Krylov iterative methods

- use of hybrids methods GMRES/MERAM_LS tuned to use two or more interconnected supercomputers thanks to the Large Scale Workflow Computing YML formalism. YML is a research project that aims at providing tools for using global computing middleware such as GRID, Peer to Peer and metacomputing environment

A*STAR Computational Resource Centre

Centre de Calcul ROMEO
MAISON DE LA SIMULATION
de Champagne-Ardenne

Perspectives

- democratise InfiniBand
- Singapore national InfiniBand fabric
- permanent InfiniBand trans-continental interconnections
- new partners and new applications

The Fusion experiment for which large ECEI fast camera data from Singapore to USA must be transferred

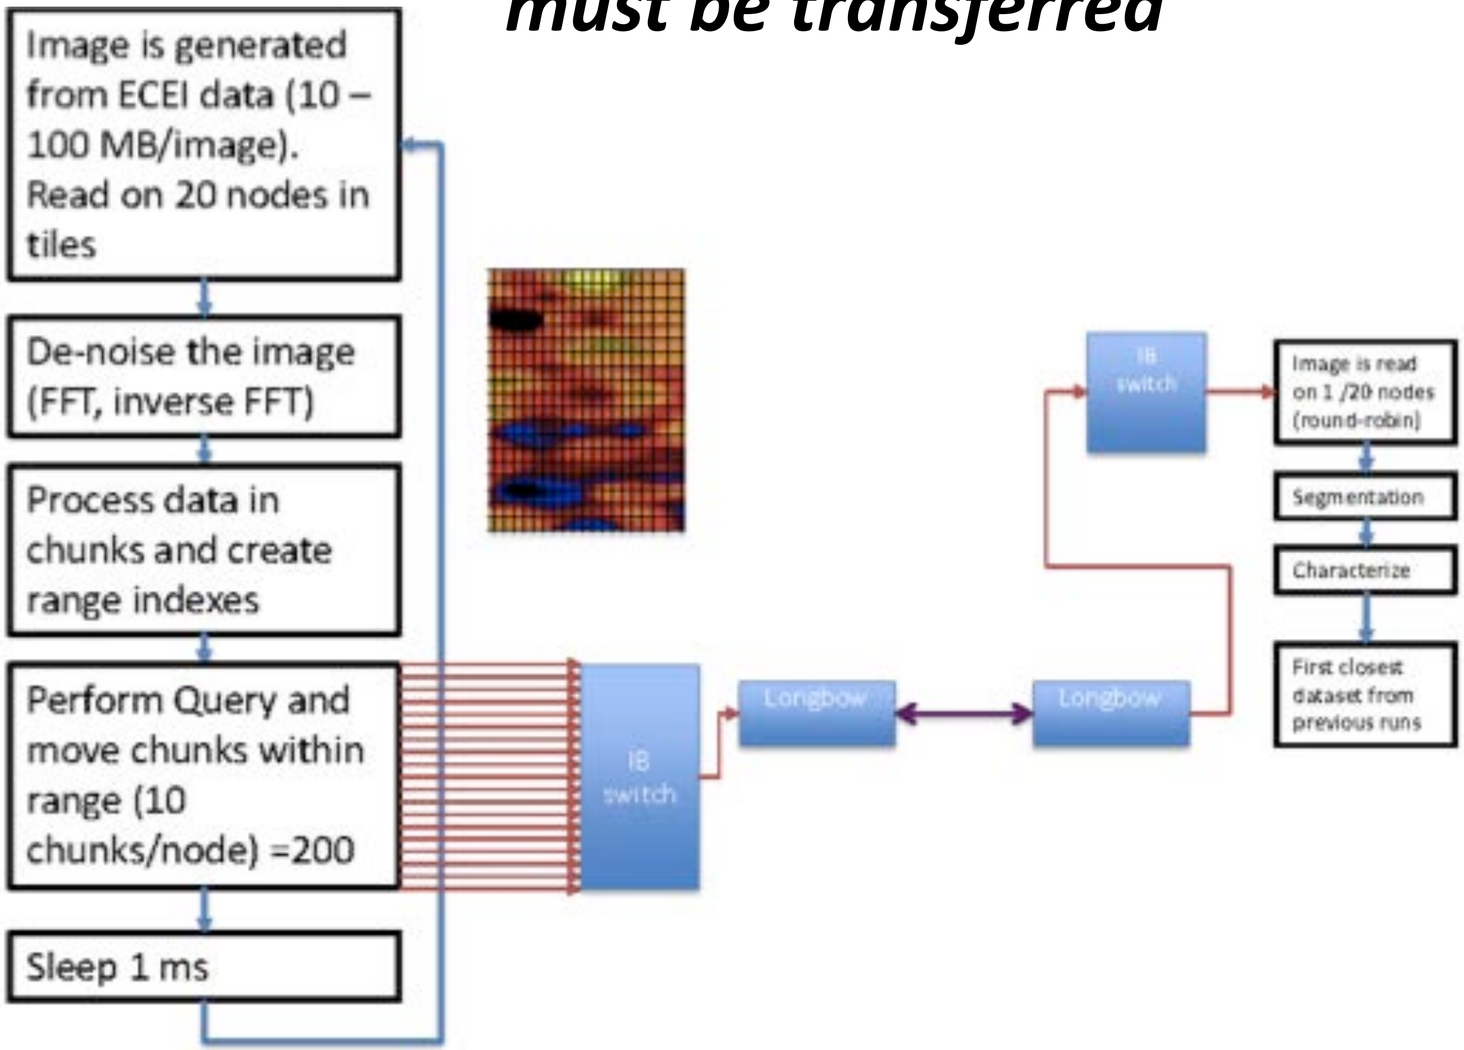

Participants: ORNL, U Tennessee, Rutgers U, A*CRC

The InfiniCortex Application Layer

Molecular Dynamics – Design of New Materials

- This demonstration is based off a scenario from materials scientists interested in understanding fracture in nano-structured materials
- It uses LAMMPS to simulate the block of nano-structured metal while under stress.
- Simulation proceeds until the first plastic deformation (start of fracture) is detected.
- At that first fracture, the system is fully characterized to understand where and, hopefully, why things broke.

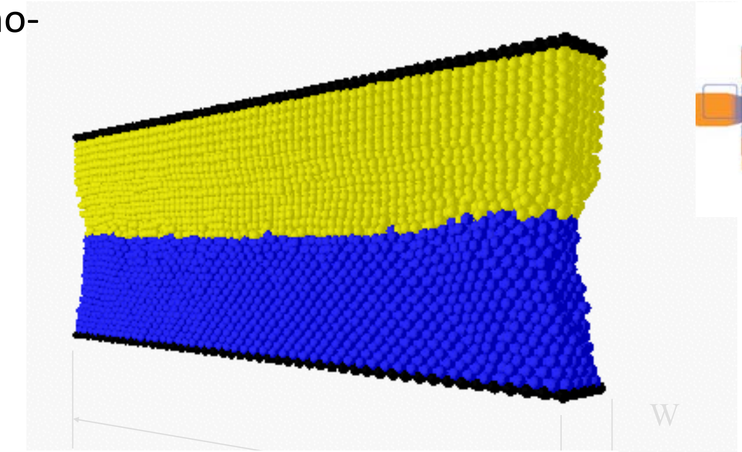

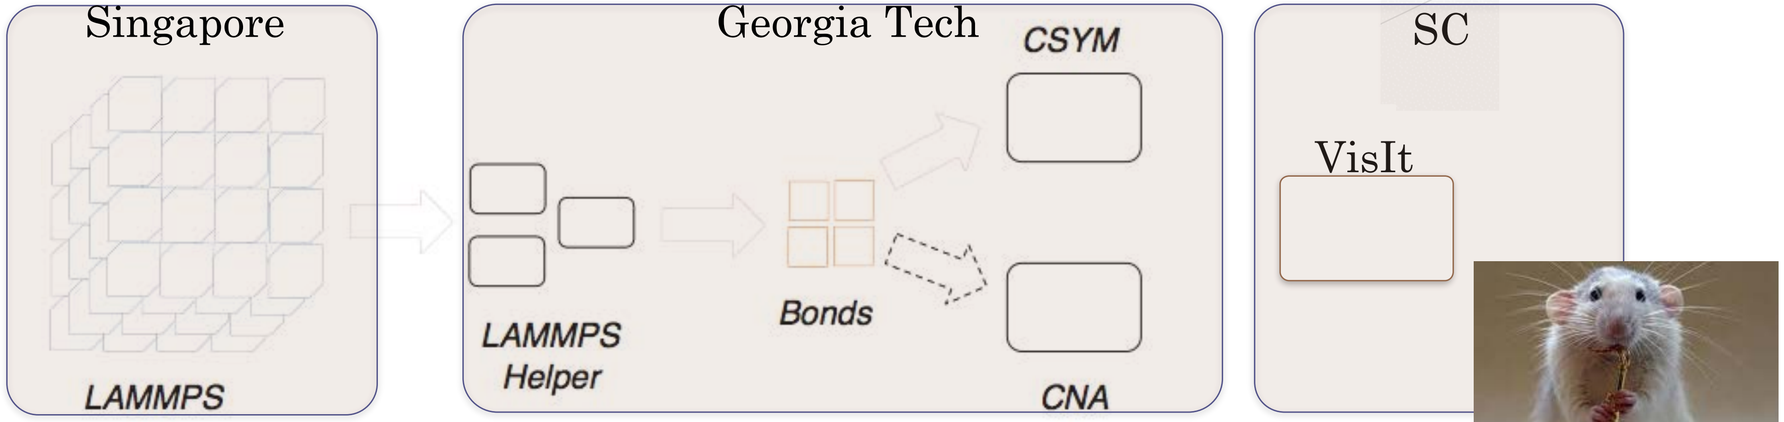

Singaporean Team

Dr Marek Michalewicz (PI)

A/Prof Tan Tin Wee (PI)

Prof Yuefan Deng

Yves Poppe,
International connectivity

A*CRC

Tan Geok Lian (Networking)
Lim Seng (Networking)
Dr Jonathan Low (H/W, S/W, Applications)
Dr Dominic Chien (S/W, Applications)
Dr Liou Sing-Wu (S/W, Applications)
Dr Gabriel Noaje (S/W, Applications)
Paul Hiew (H/W)

With help from:

SingAREN
- **A/Prof Francis Lee**
- **Prof Lawrence Wong**

NTU
- **Stanley Goh**

A*STAR Computational Resource Centre

American Team

Oak Ridge National Lab
Dr Scott Klasky
Dr Jong Choi

Georgia Tech
Prof Matthew Wolf
Prof Greg Eisenhauer

Rutgers University
Prof Manish Parashar

Oak Ridge National Lab
Dr Scott Klasky

Stony Brook University
Prof Deng Yuefan
Prof Tahsin Kurc

University of Tennessee
Glenn Brook

Stony Brook University
Prof Deng Yuefan

Australian Team

Canberra: National Computing Infrastructure

Andrew Howard, NCI

Jakub Chrzeszczyk, ANU and Kenneth Ban NUH/NUS/A*CRC

Obsidian Strategics

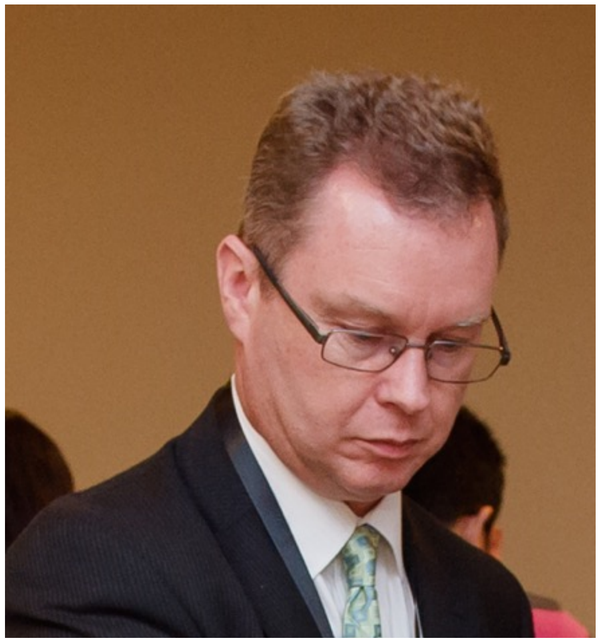

Dr David Southwell, CVO

Jason Gunthorpe, CTO

Team Singapore

Singapore

Team Singapore

Singapore

Team Australia

Canberra

The Japanese Team

TiTech: Tsubame-KFC: Prof Satoshi Matsuoka, S. Muira *et al.*

TEIN-JP NOC Team, KDDI

SINET Team, NII

NICT

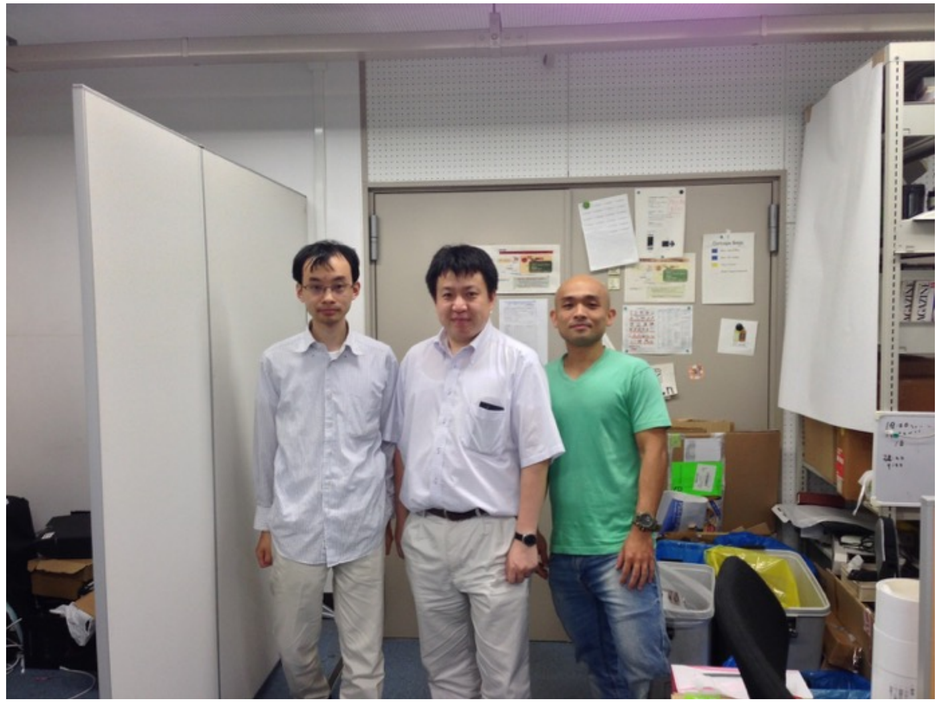

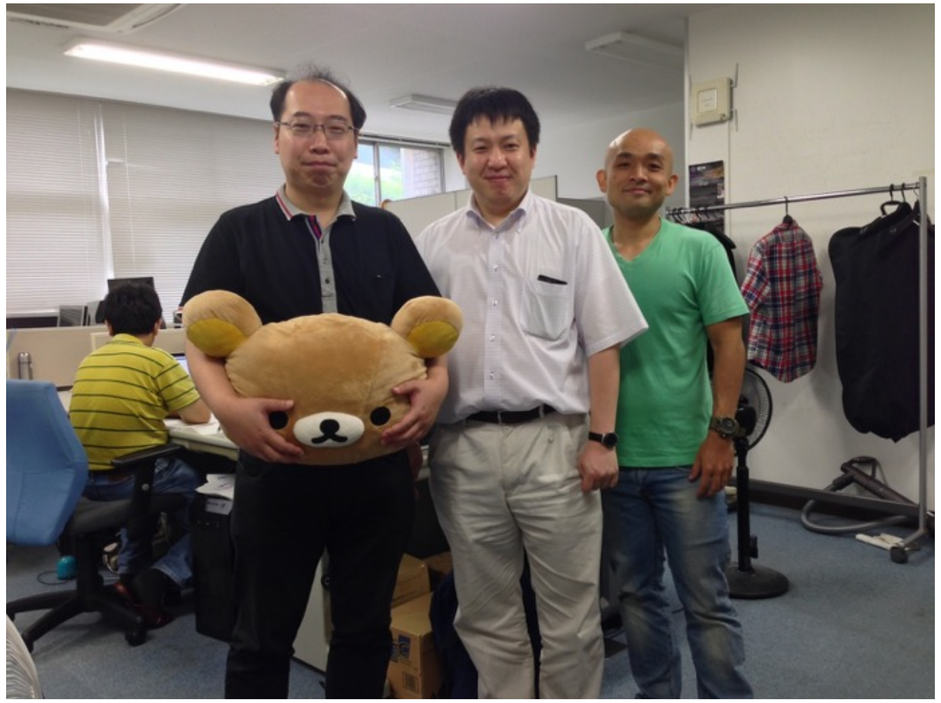

Japan

System Software: BGFC - Next Generation InfiniBand Subnet Manager

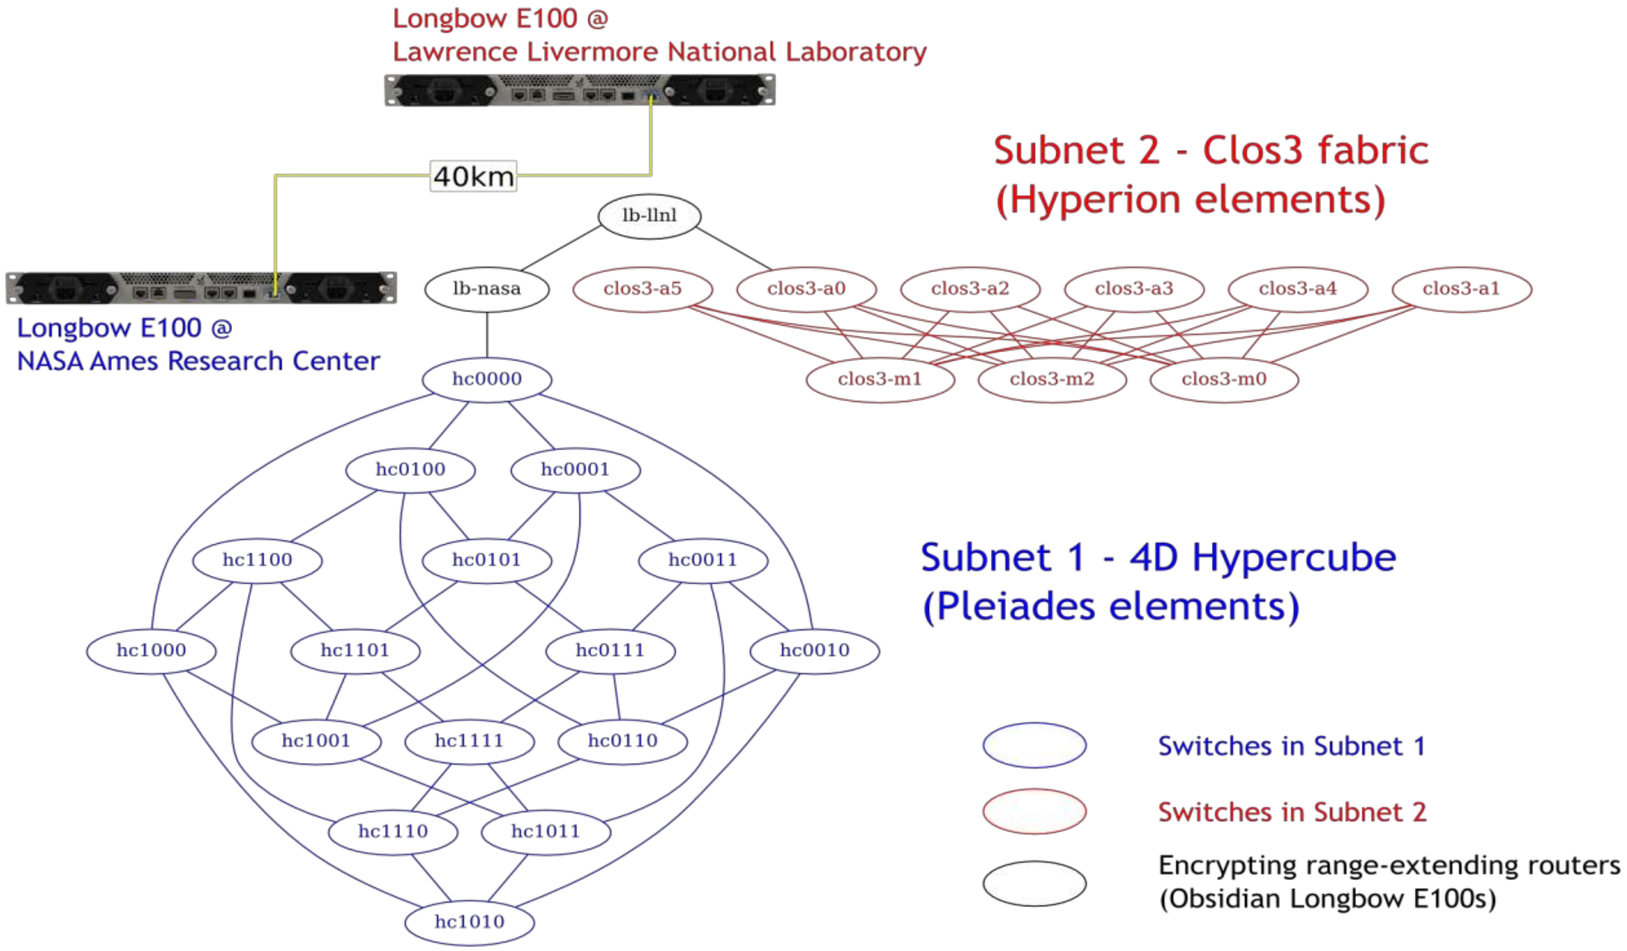

A 1Gbyte/second 40km routed and encrypted InfiniBand link was demonstrated at SC12, using BGFC to orchestrate the fabric comprising two subnets with very different internal topologies (“compound topology”).

OBSIDIAN STRATEGICS

Computational Resource Centre

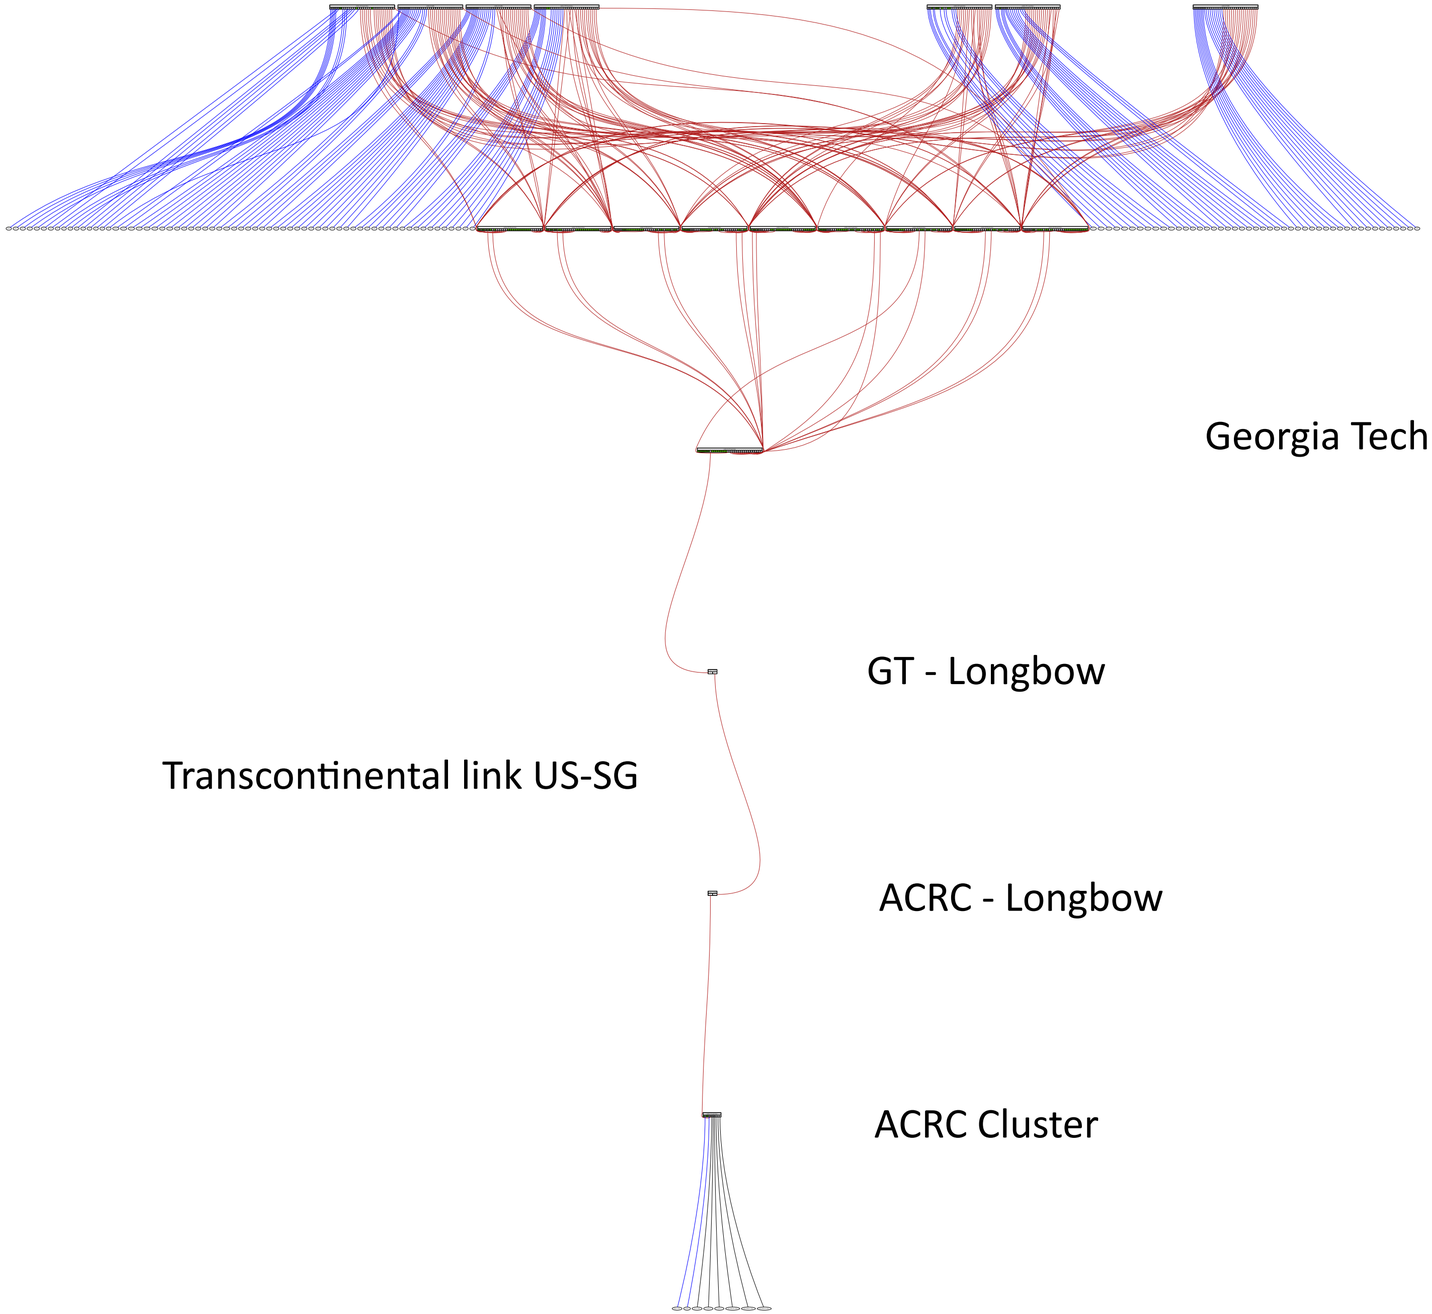
Georgia Tech
GT - Longbow
Transcontinental link US-SG
ACRC - Longbow
ACRC Cluster

SG to AU Observed data transfer rate at 10Gb using Obsidian DSYNC

1 second	900 MB
1 minute	54 GB
1 hour	3.24 TB
1 day	77 TB
1 week	539 TB
1 month	2.3 PB