



BDEC Collaboration Opportunities at A*STAR and Singapore

A*STAR Computational Resource Centre

Marek Michalewicz

4th BDEC Workshop @ISC15

Frankfurt 16 July 2015

The Big Picture:

Singapore's National Supercomputer Centre

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

- National Supercomputing Centre (NSCC)
 - New 1+ PetaFLOP Supercomputer
 - Recurrent investment every 3-5 years
 - Co-investment (OpEx) from primary stakeholders
- Science, Technology and Research Network (STAR-N)
 - A high bandwidth network to connect the distributed login nodes
 - Provide high speed access to users (both public and private) **anywhere**
 - Support transfer of large data-sets (both locally and internationally)
 - Builds **local and international network connectivity** (Internet 2, TEIN*CC)
 - ASEAN, USA Europe, Australia, Japan, Middle East

Funding (MTI) and co-funding (A*STAR, NUS, NTU) approved Nov. 2014

Tender Open: 20th January 2015

Tender Closed: 14th April

Tender Awarded: 15th June 2015

Facility open to users: 1st week October 2015

| Document No. | Tenders and Quotations | Publication Date | Closing Date / Time | Status |
|---|--|------------------|---|---------|
| Tender: NST000ETT15000015 Tender Ref: 704/2/2-87 | Supply, Installation & Maintenance of a Petascale Supercomputer System & Managed Services, including A&A of Data Centre Suitable to Host the System for the NSCC Calling Entity Agency for Science, Technology and Research Award Date: 15/06/2015 | 20/01/2015 | 14/04/2015 4:00 PM Electronic Submission | Awarded |



Login (with)
 ✦ SingPass
 ✦ GeBIZ ID

[Home](#) >> [Past Opportunities](#) >> [Tenders and Quotations](#) >> [Award](#)



TENDER NOTICE

Notice | Schedule | **Award**
 HEADER :: **DETAILS** Back to Listing

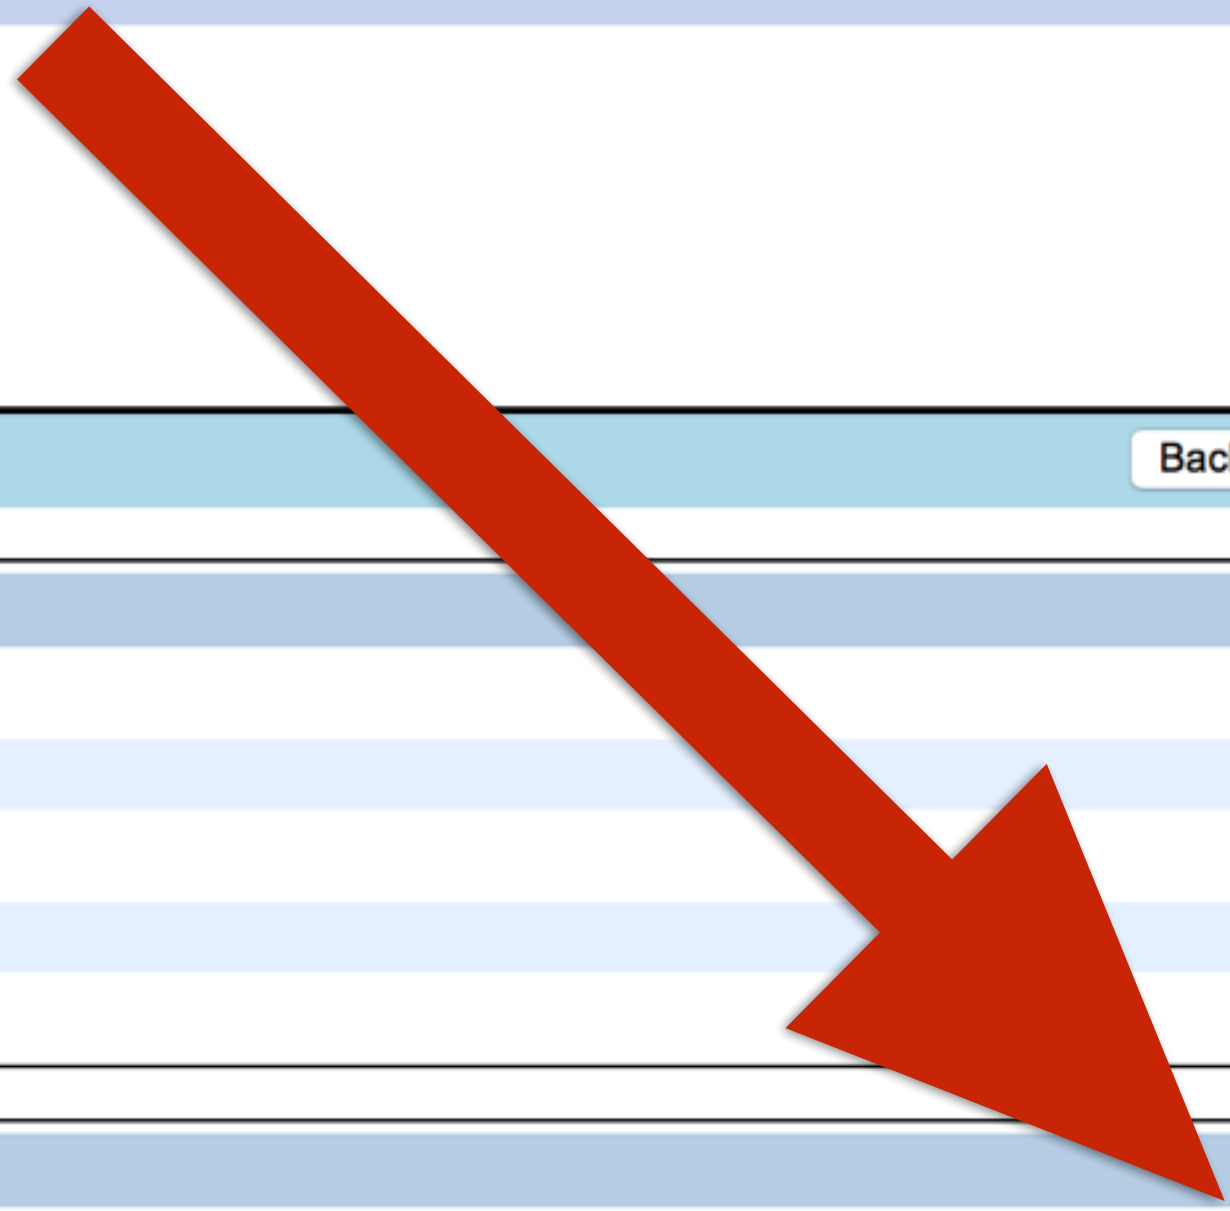
| AWARDED SUPPLIER 1 | |
|-------------------------|--|
| Name of Supplier | FUJITSU ASIA PTE LTD |
| Address | Nexus @ one north 1 Fusionopolis Link, #04-01 Singapore 138542 |
| Awarded Value | SGD 45,241,669.87 |
| Consortium | No |
| No. of Companies | 0 |

| Awarded Items | | | | | | |
|---------------|--|--------|-----|----------|-----------------|---------------|
| S/No. | Description: | Qty | UOM | Currency | Unit Price | Awarded Value |
| 1. | Petascale Supercomputer System (Hardware & Software) | 1.0000 | LT | SGD | 36,624,473.8700 | 36,624,473.87 |
| 2. | Addition and Alteration Works of Data Centre | 1.0000 | LT | SGD | 5,611,520.0000 | 5,611,520.00 |
| 3. | Manage Services - 3 years | 1.0000 | LT | SGD | 3,005,676.0000 | 3,005,676.00 |

Business Alerts
 ✦ How to Subscribe?

Business Opportunities
 ✦ Tenders & Quotations
 ✦ Qualifications
 ✦ Auctions
 ✦ Request for Information

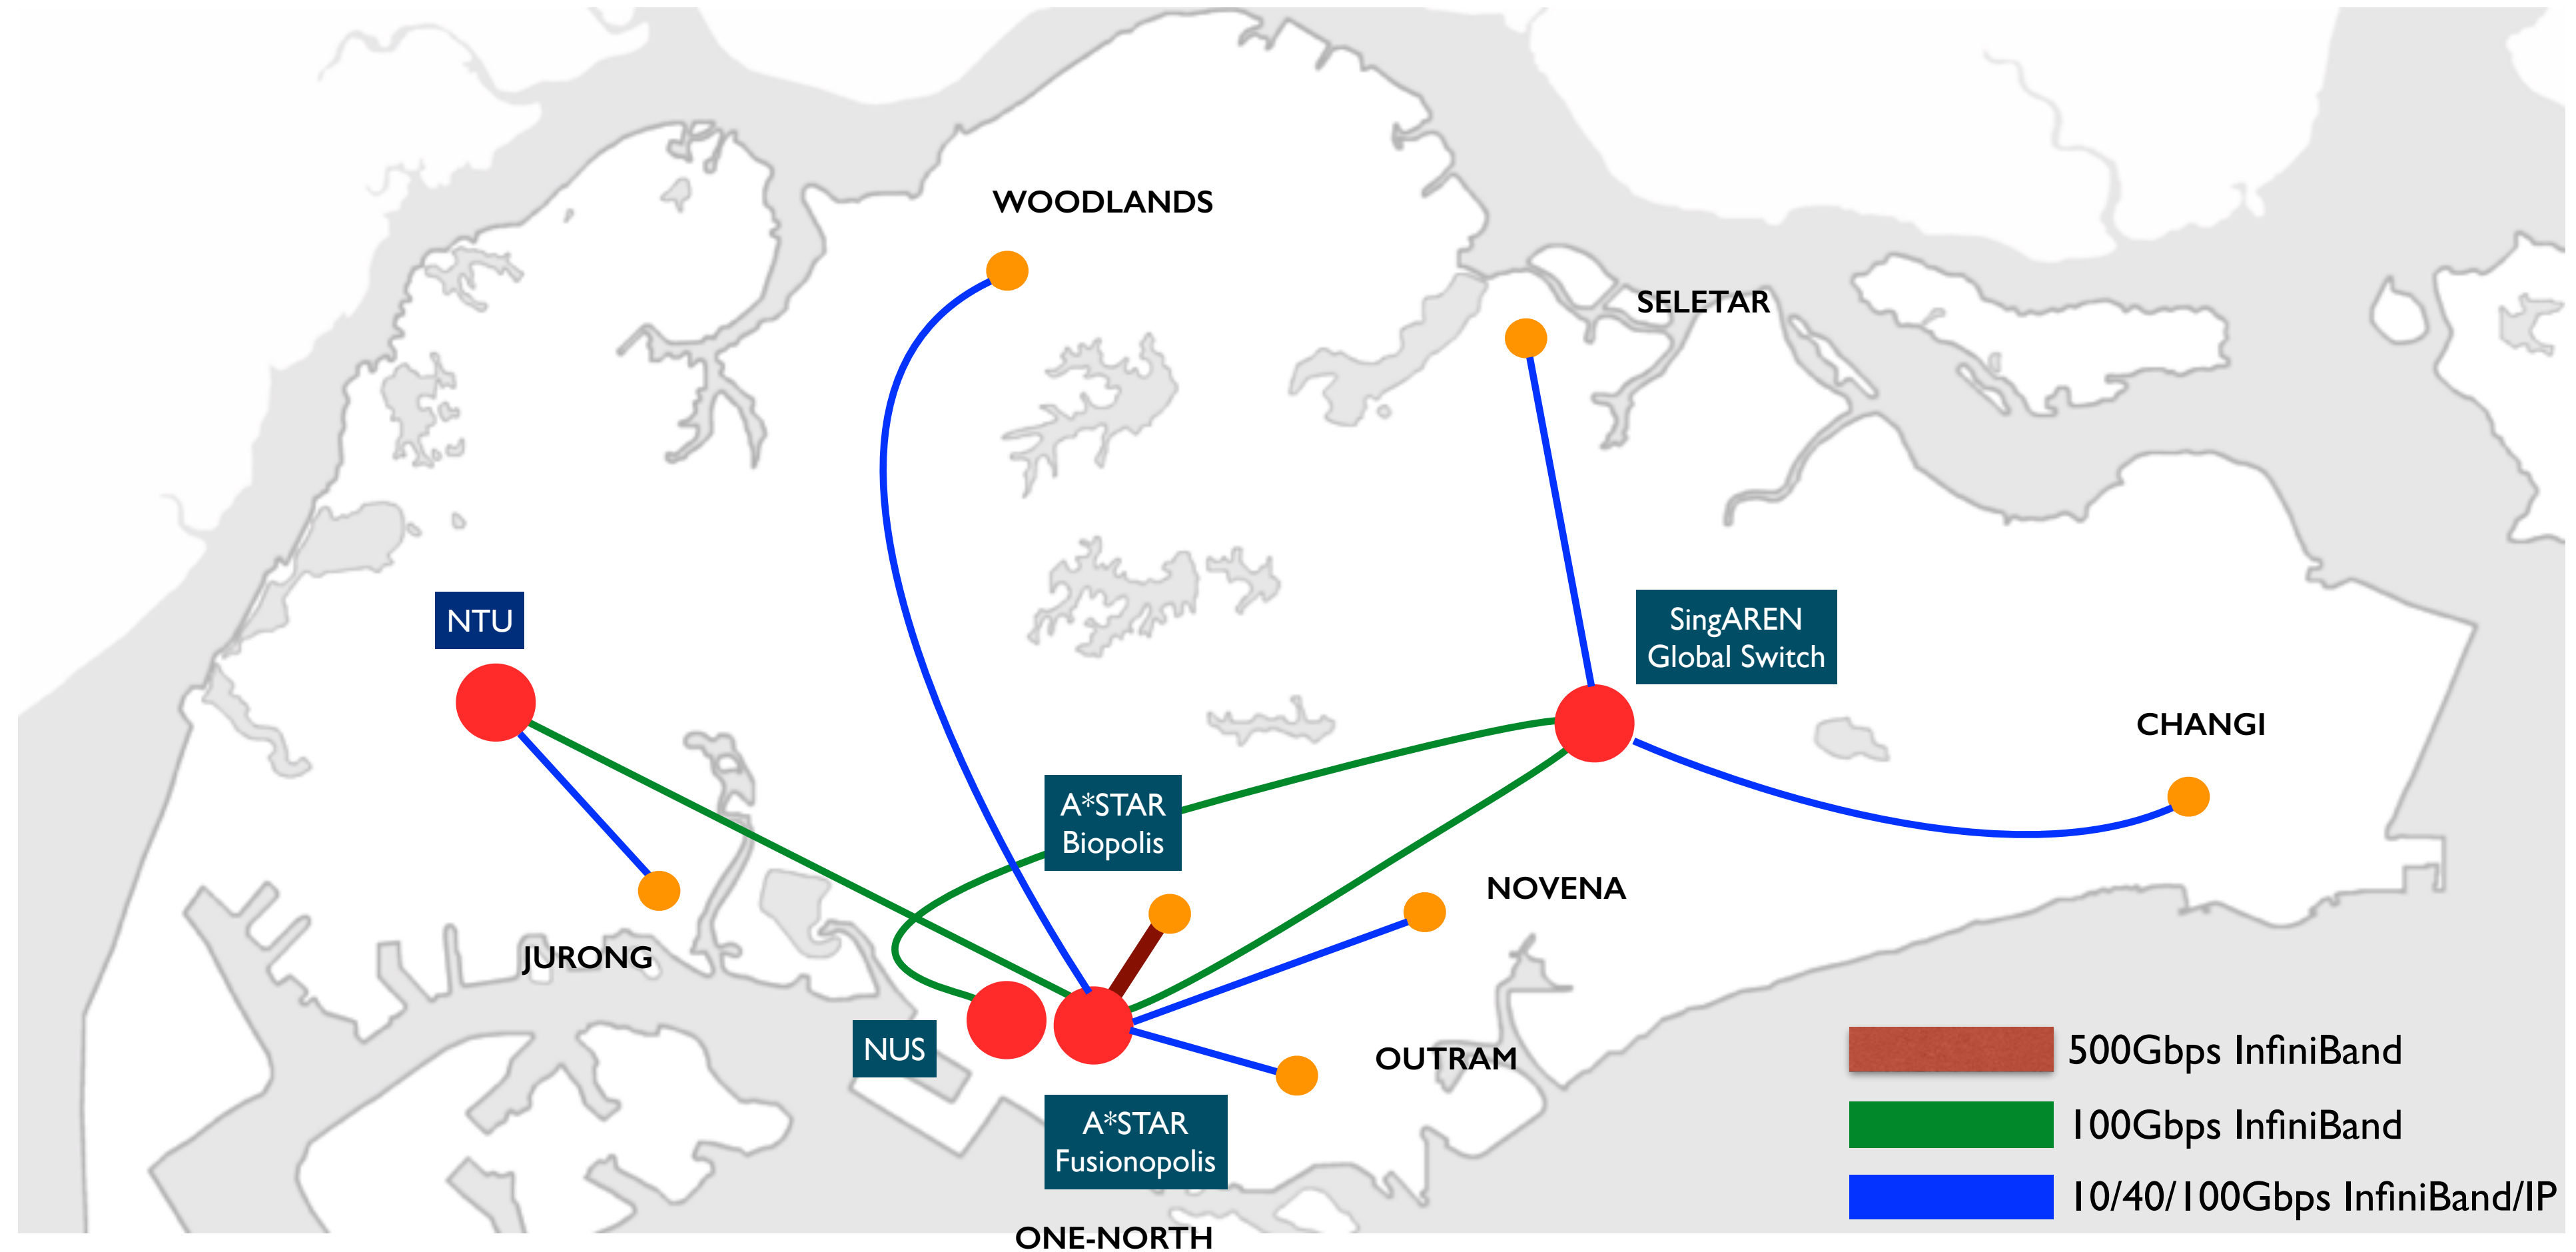
Past Opportunities
 ✦ **Tenders & Quotations**
 ✦ Qualifications
 ✦ Auctions
 ✦ Request for Information



Some features of NSCC Supercomputer

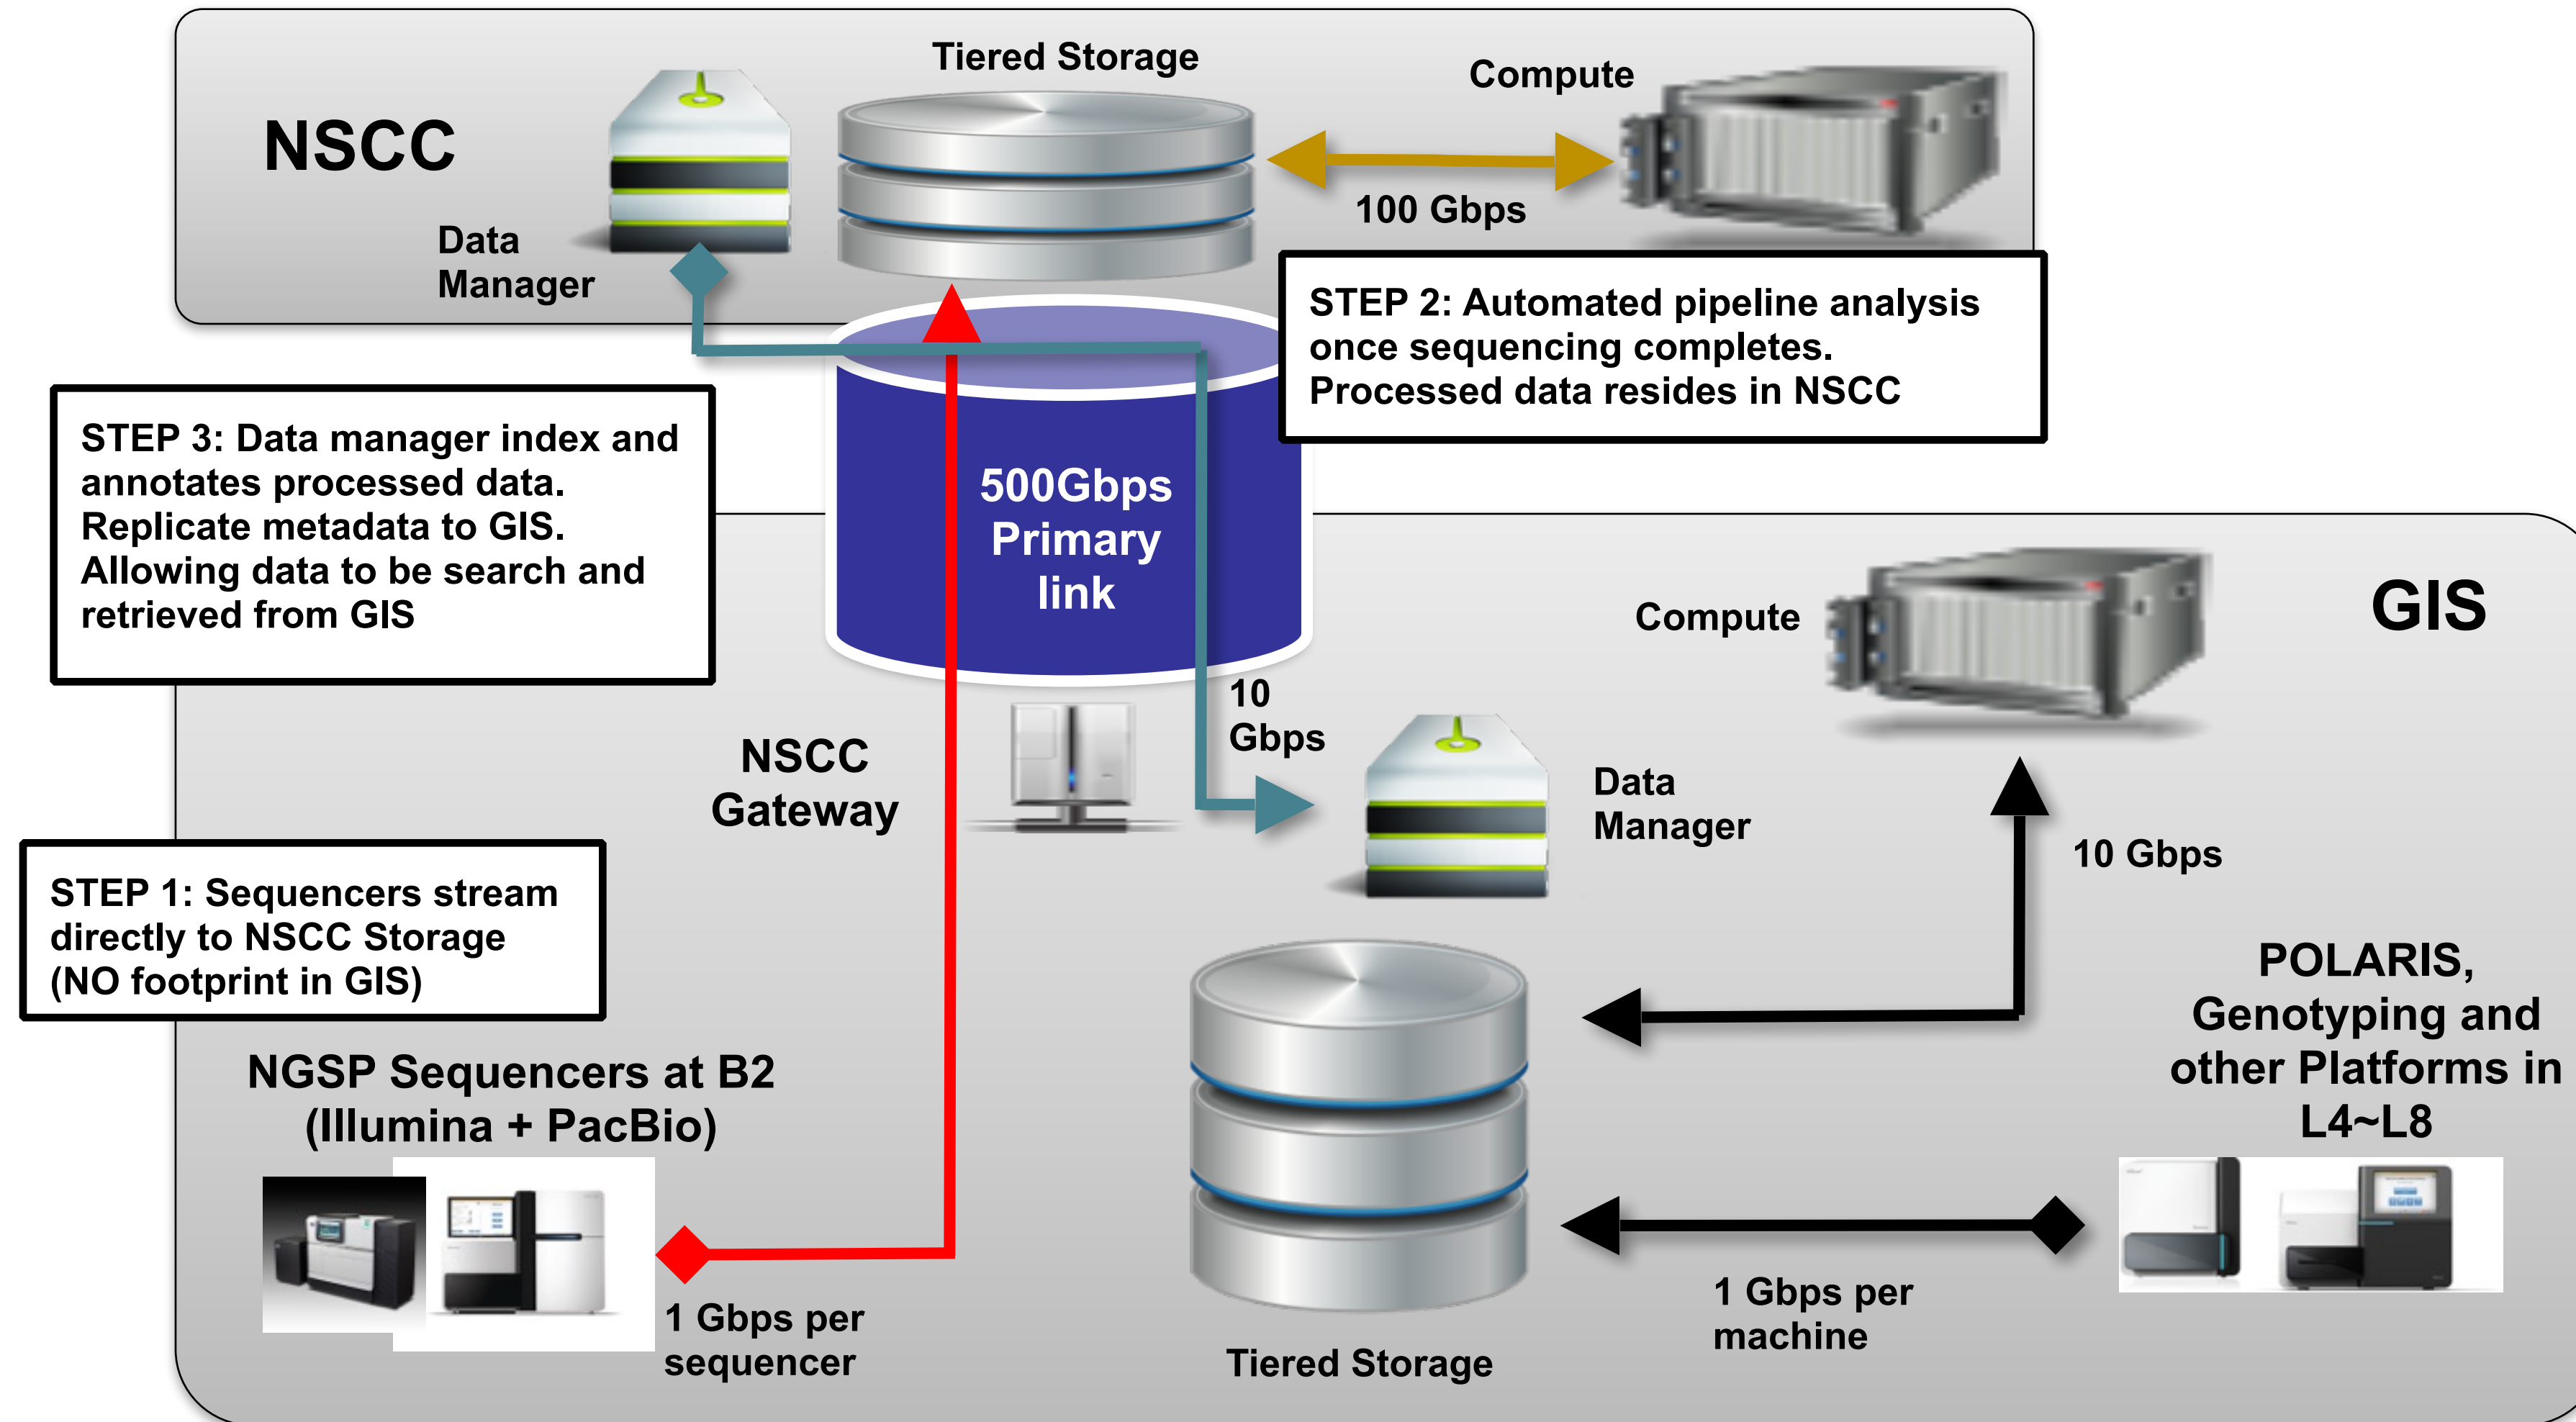
1. Base system: ~1 PFLOPS: 1,288 nodes, dual socket, 12 cores per CPU E5-2690v3 (Fujitsu)
2. Mellanox EDR Fat Tree
3. Storage: **10PB**, HSM, Tiered, Lustre + GPFS + WOS
4. **I/O (minimum) 500 Gbps flash burst buffer** , 10x Infinity Memory Engines (IME), (DDN)
5. Application Software - National license options - scheduler, Allinea tools, NAG
6. **10 Large memory nodes**: 1x 6TBytes, 4x 2TBytes, 5x 1TBytes
7. Workflow pipeline (from sequencer to memory/storage) + **interactive access** built in
8. **500 Gbps pipeline** between Biopolis and Fusionopolis for genomics workflows
9. InfiniBand connection to all end-points (login nodes) at three University campuses and industrial sites
10. Just-in-time resource delivery i.e. **interactive access** for some workloads (e.g. genomics)
11. 128 (10%) nodes GPU accelerated: Tesla K40 (NVIDIA)
12. Very highly instrumented (time, energy, power): high frequency measurements at component, chassis and cabinet level
13. Warm water cooled

Singapore InfiniBand connectivity



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

GIS-NSCC Integration: Future



InfiniCortex

***creating the Global Supercomputer
and enabling research collaborations***



InfiniCortex is ...


NOT GRID!

NOT CLOUD!

NOT “Internet”!



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 Components

1. Galaxy of Supercomputers

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

2. ACA 100 & ACE 10

- Asia Connects America 100 Gbps (trial: Oct-Nov 2014, permanent October 2015)
- Asia Connects Europe: 10Gbps growing to 40Gbps October 2015

3. InfiniBand over trans-continental distances

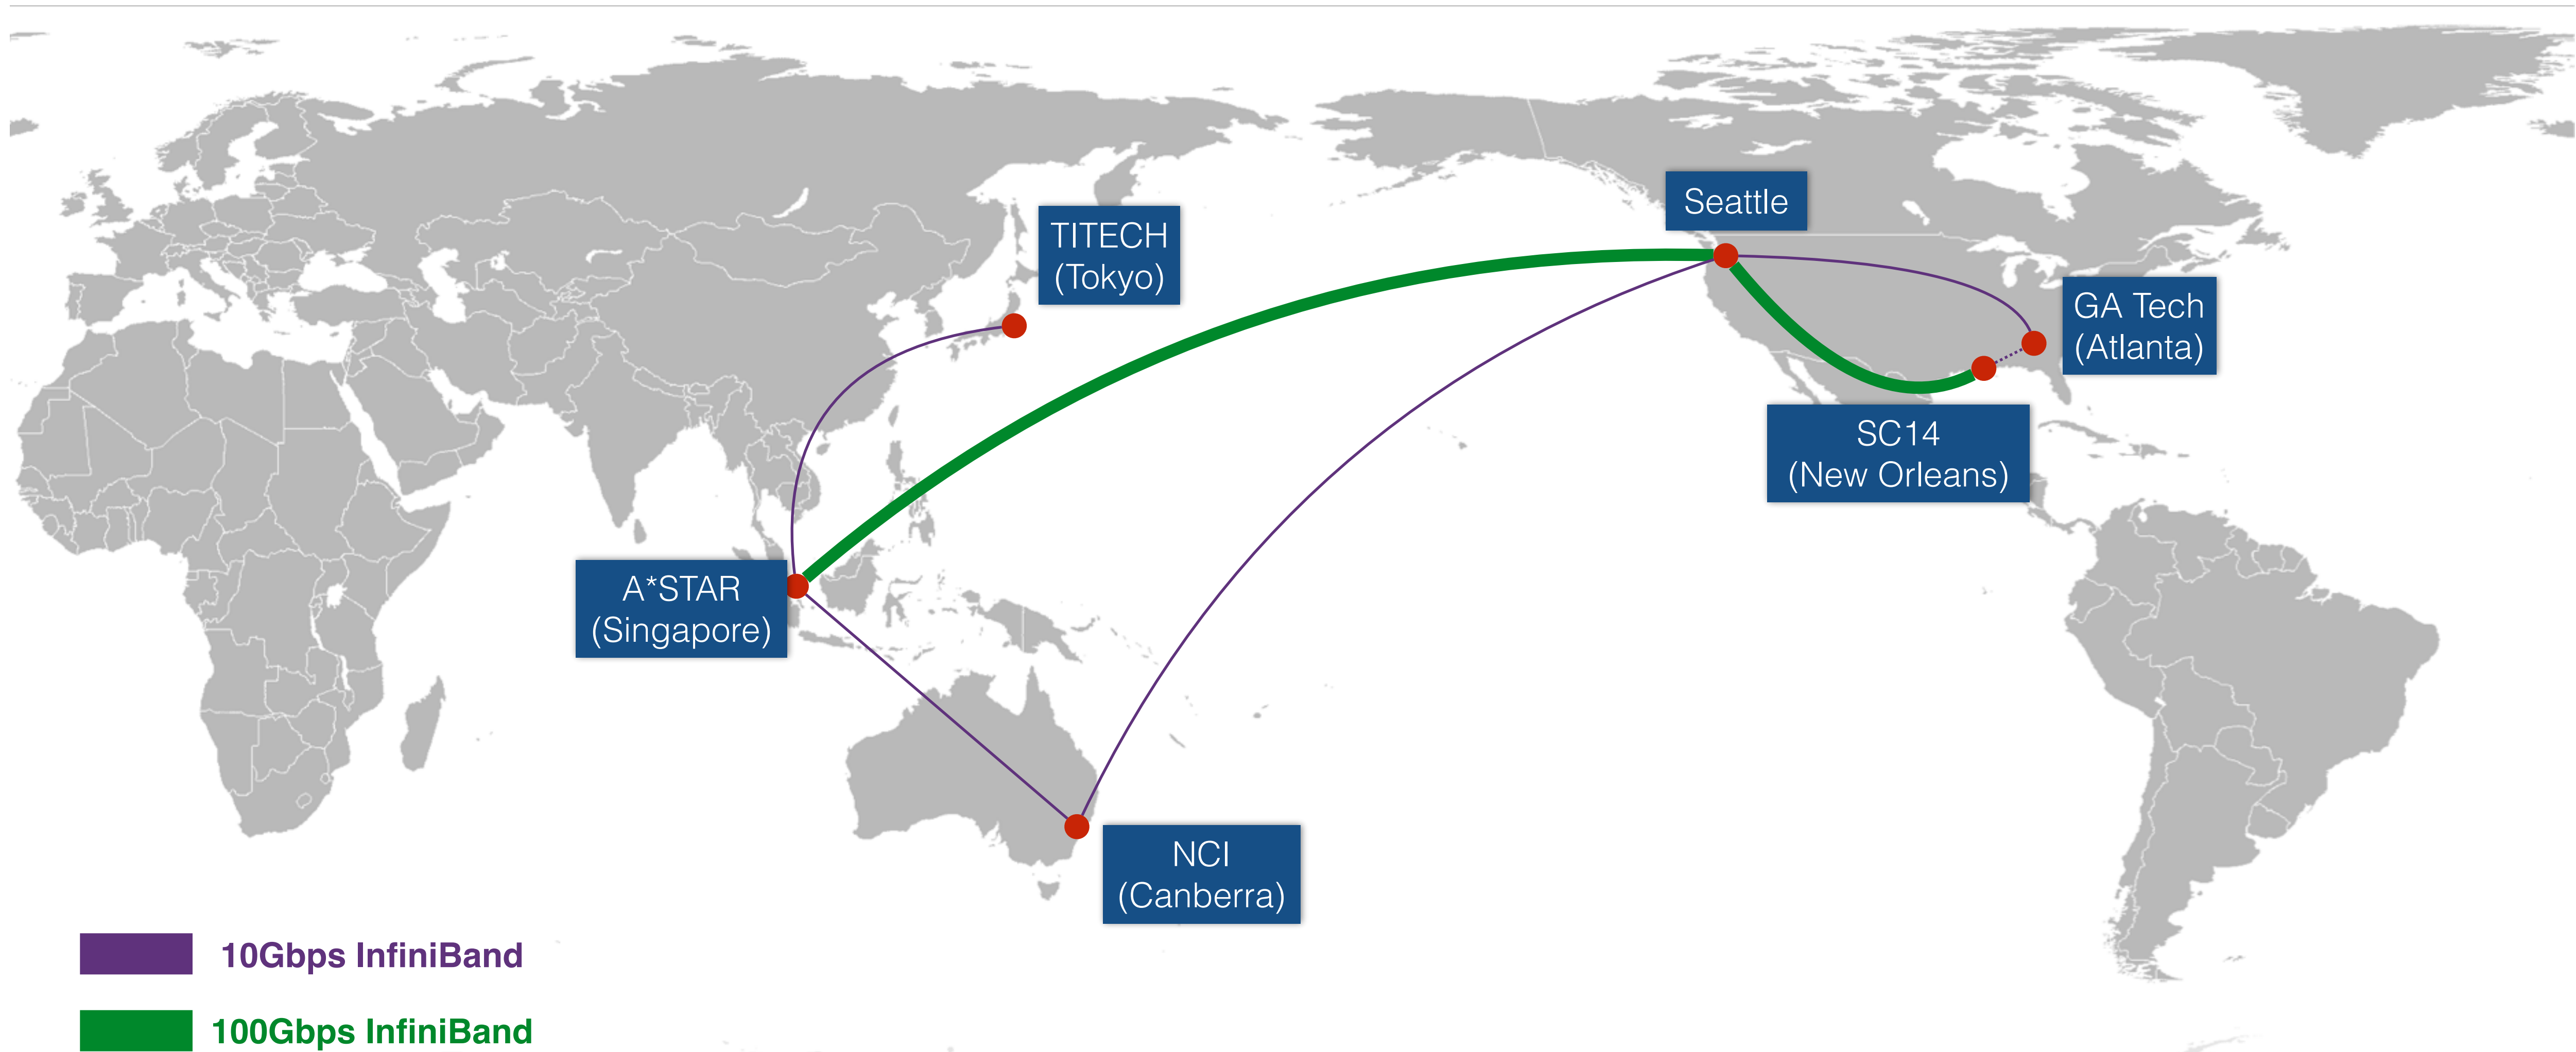
- Using Obsidian Strategics Longbow range extenders

4. Application layer

- from simplest file transfer: dsync+
- to complex workflows: **ADIOS**, multi-scale models
- **InfiniCloud**
- bare metal or virtual machine provisioning (work with Bright Computing)

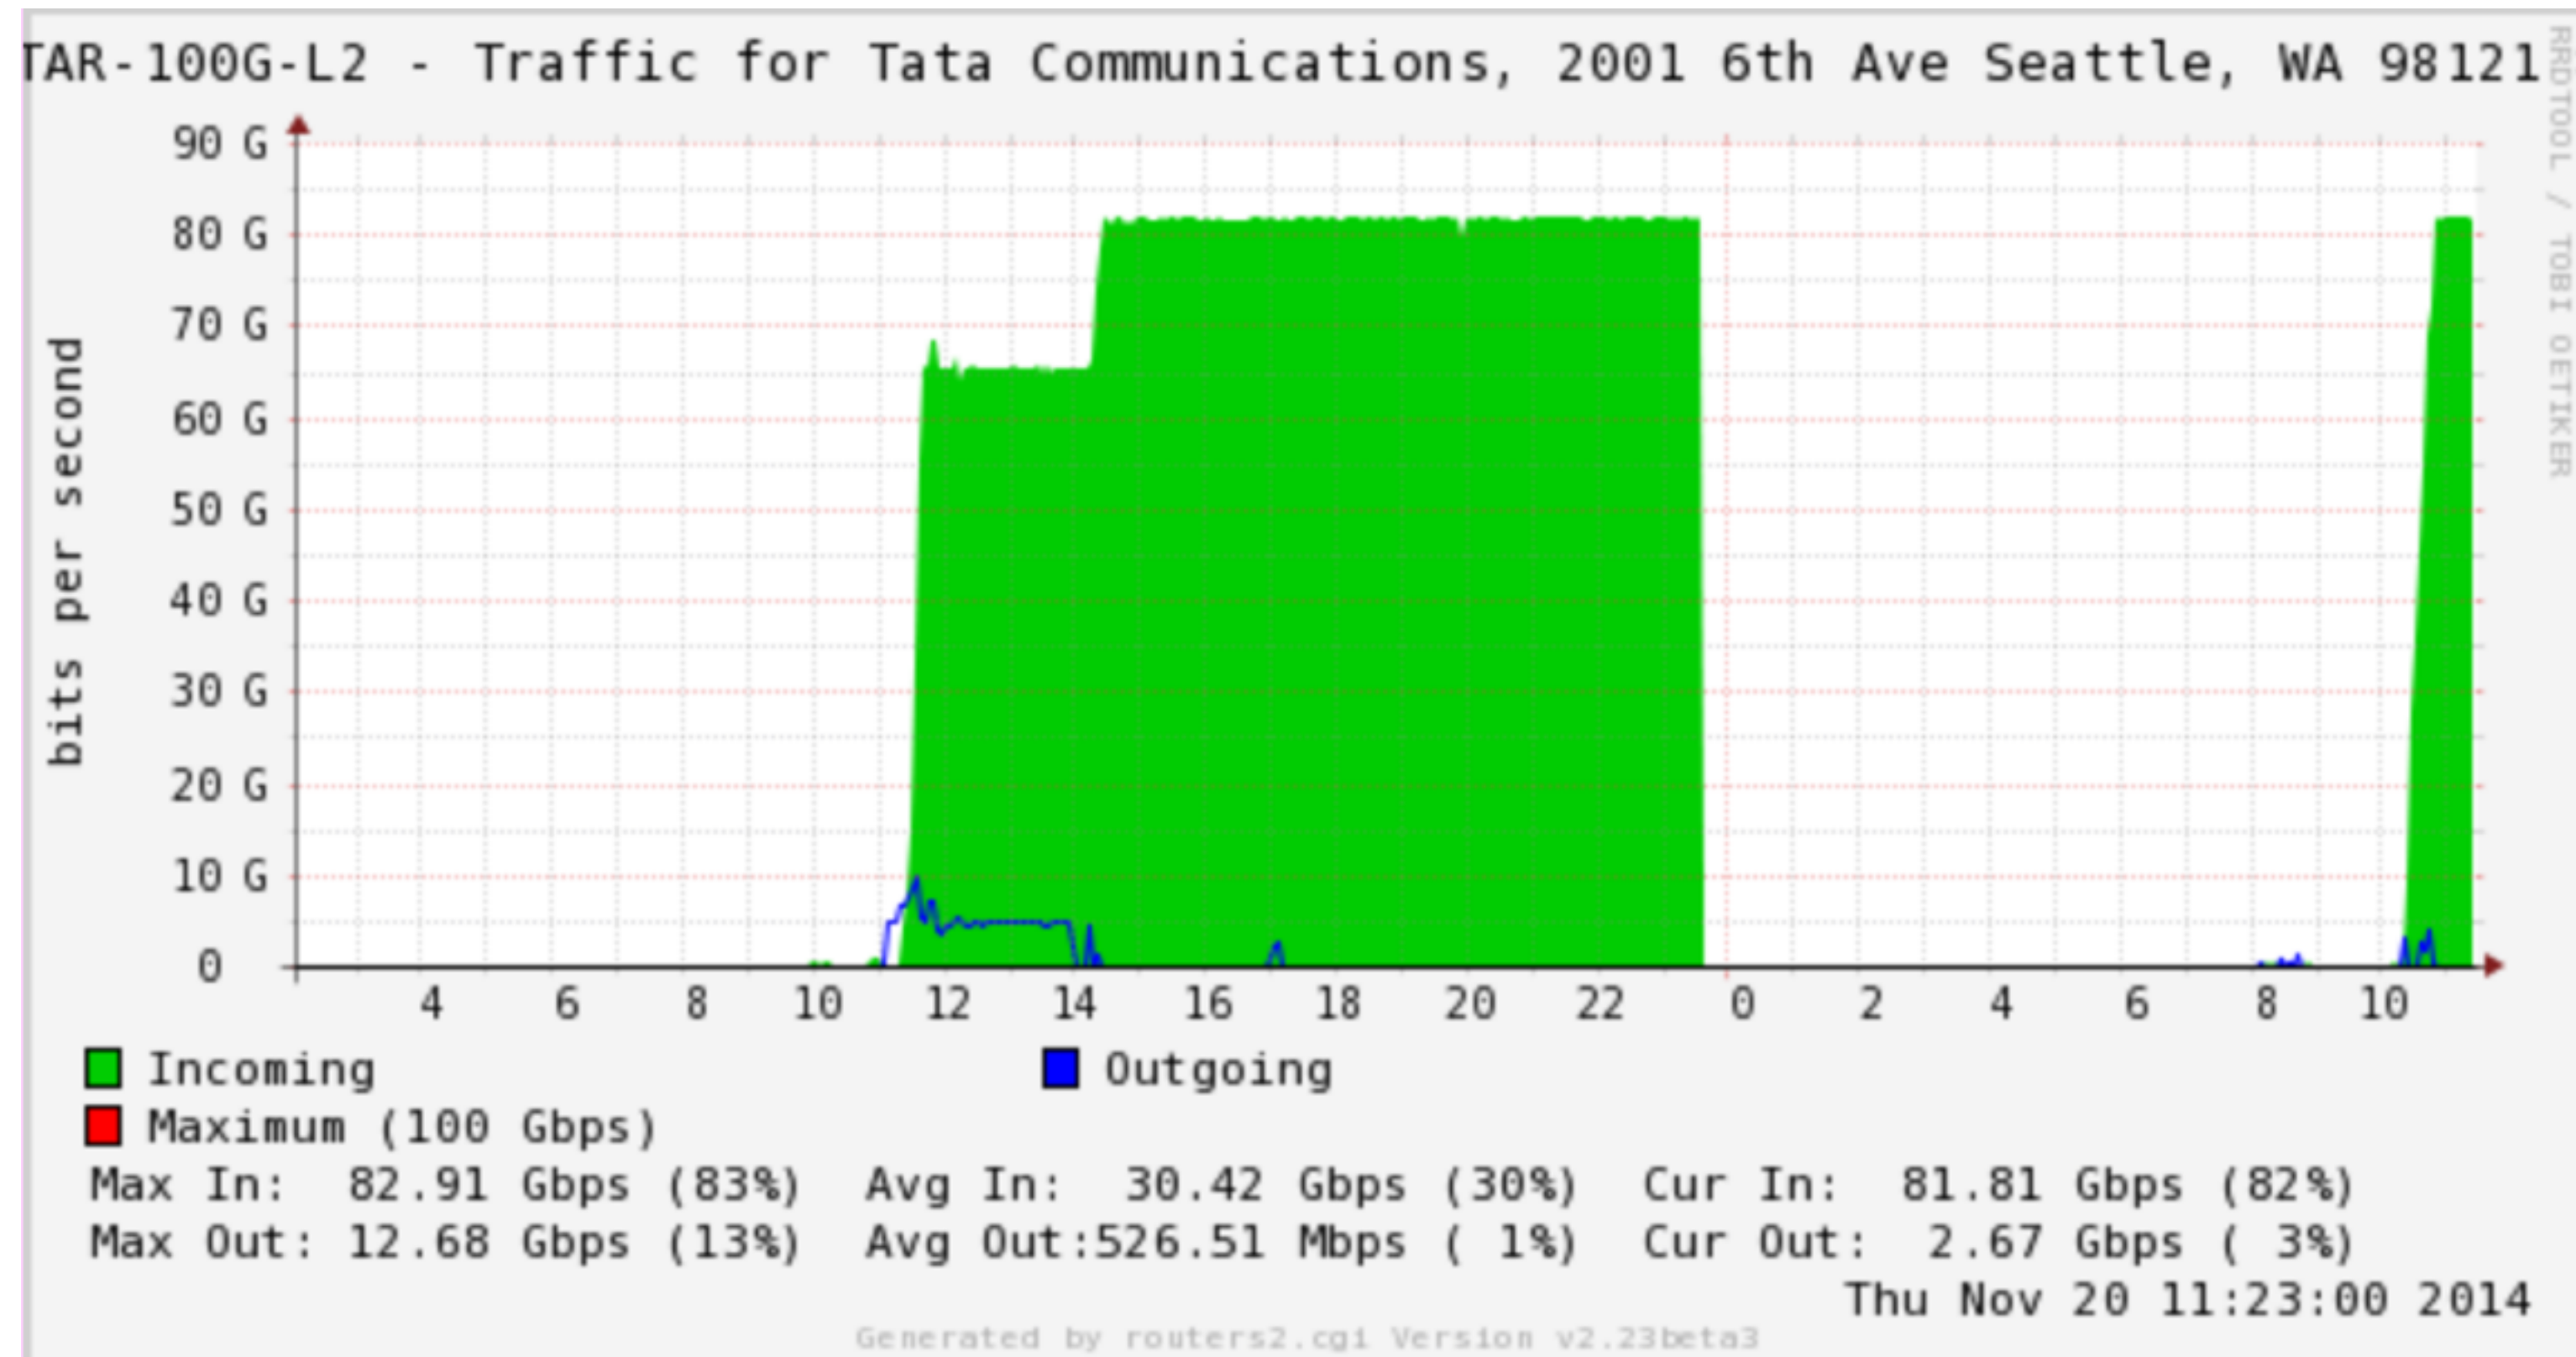


InfiniCortex 2014 (phase 1)



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.

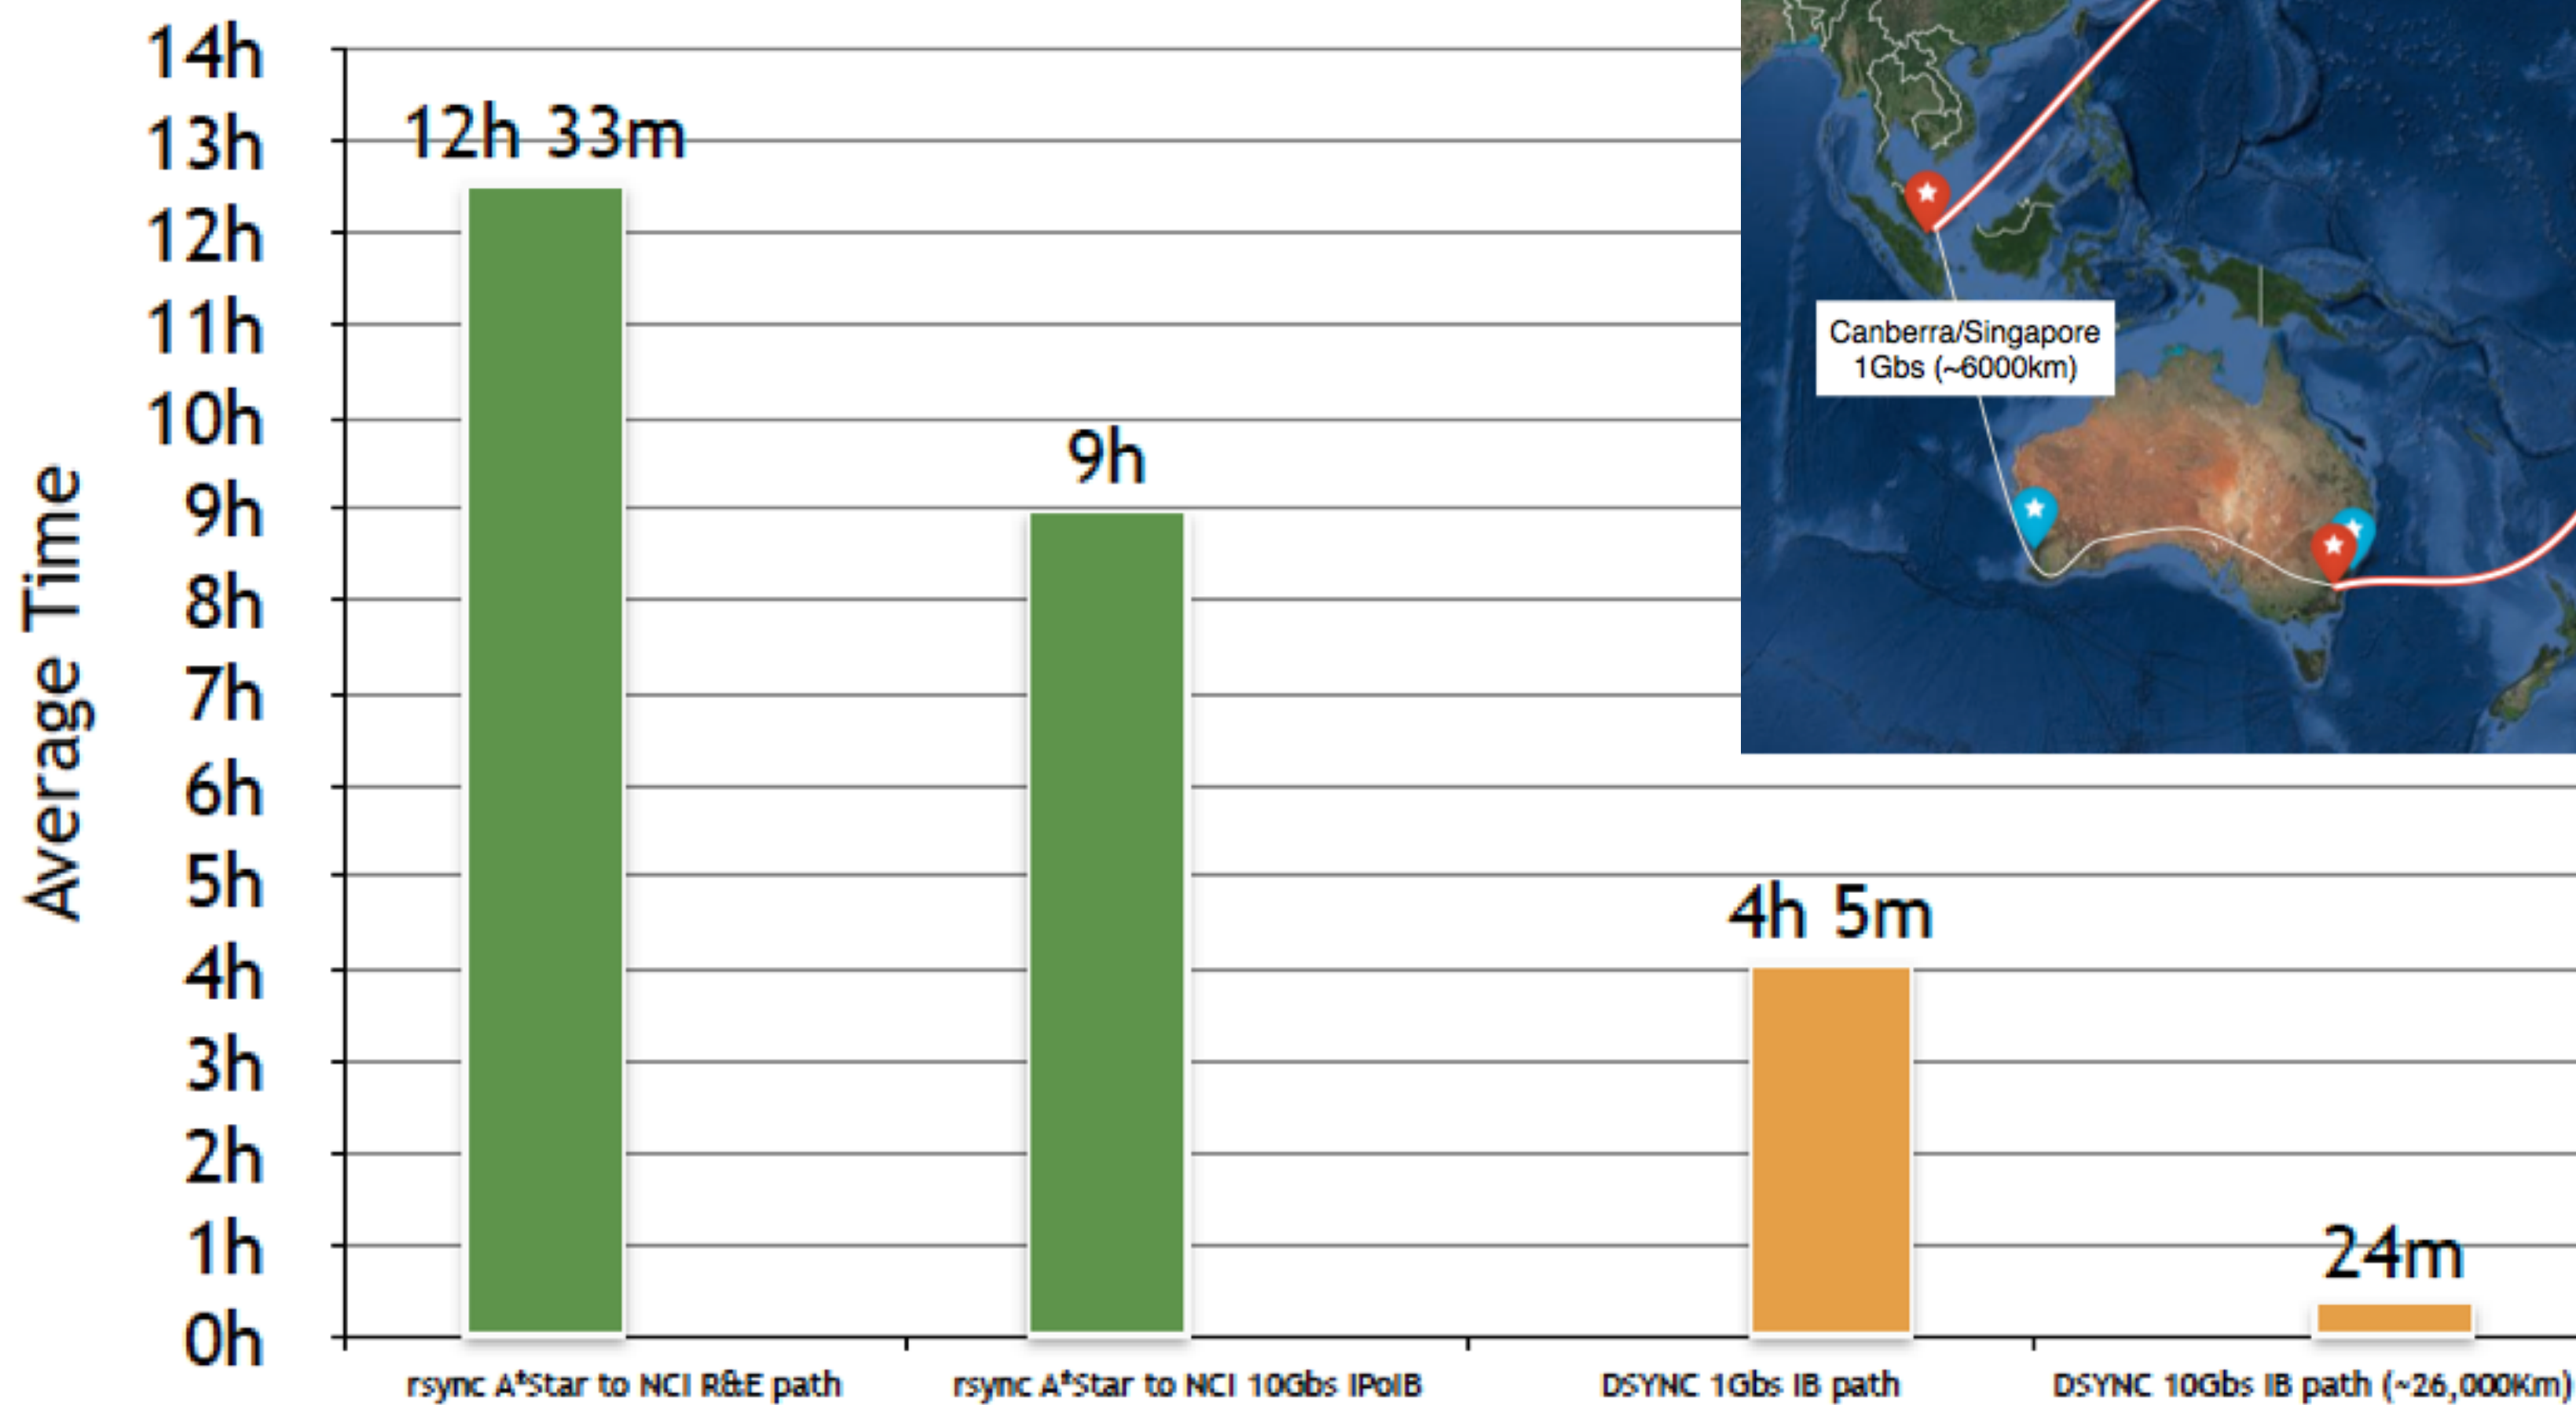
100Gbps Bandwidth Utilization





NCI

SG to AU data transfer speed



1,143Gb data set

| | |
|------------------------|-----------------------|
| <i>1 second</i> | <i>900 MB</i> |
| <i>1 minute</i> | <i>54 GB</i> |
| <i>1 hour</i> | <i>3.24 TB</i> |
| <i>1 day</i> | <i>77 TB</i> |
| <i>1 week</i> | <i>539 TB</i> |
| <i>1 month</i> | <i>2.3 PB</i> |

Project Partners

Team - Singapore



Team - USA



Team - Japan



Team - Australia



Commercial Partners



Commercial Sponsors



Awards

- Ministry of Trade and Industry (MTI) 2015 Gold Award for Innovative Project

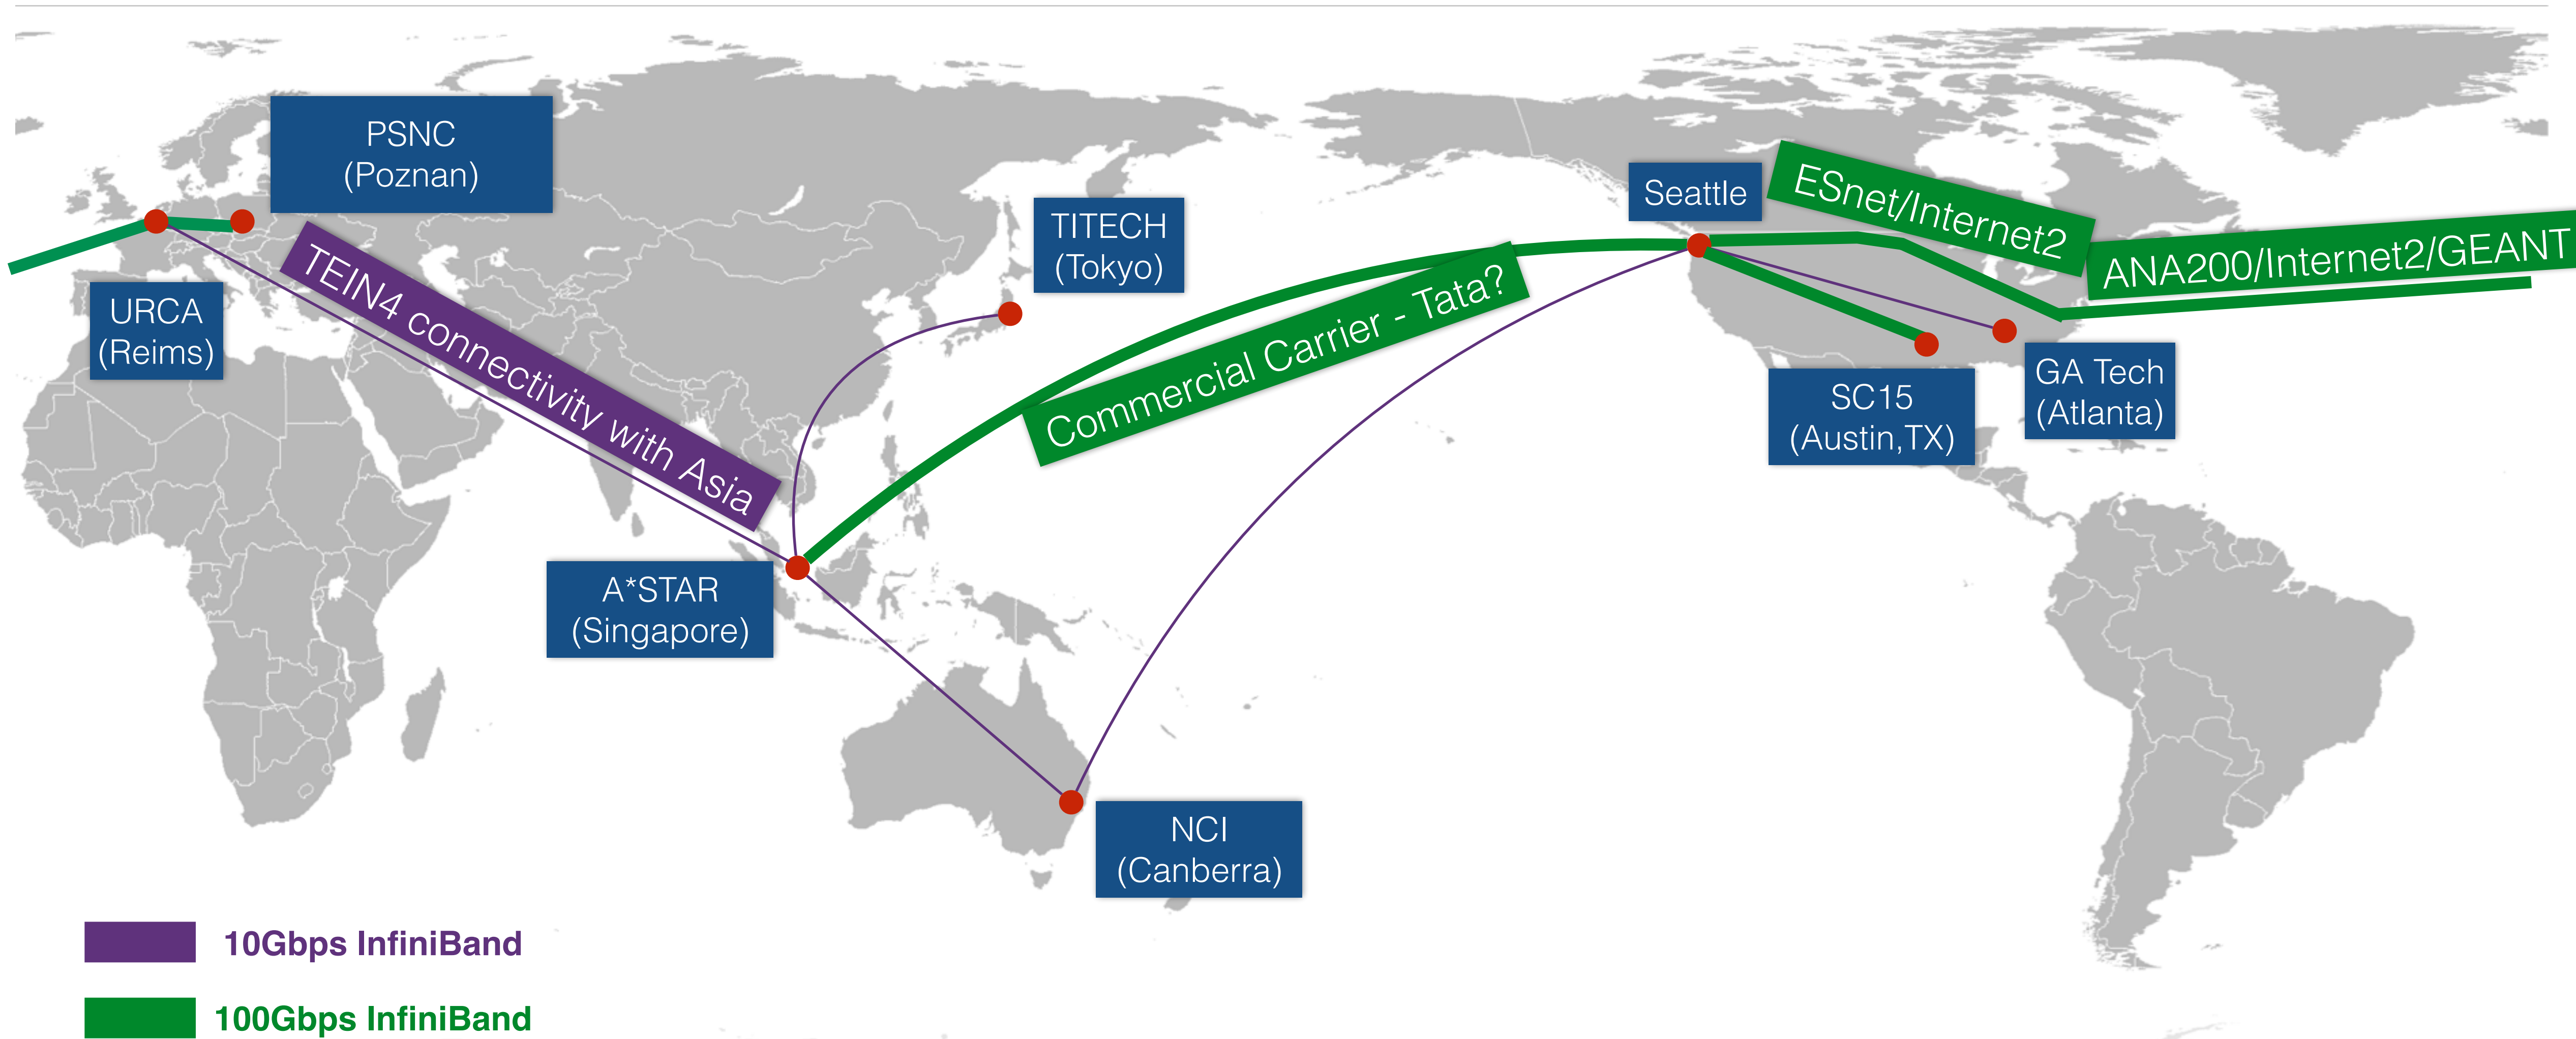


- 2015 A*STAR Innovation Award
- FutureGov Singapore Award 2015 in Technology Leadership category

- CIO 100 HONOUREE 2015



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)

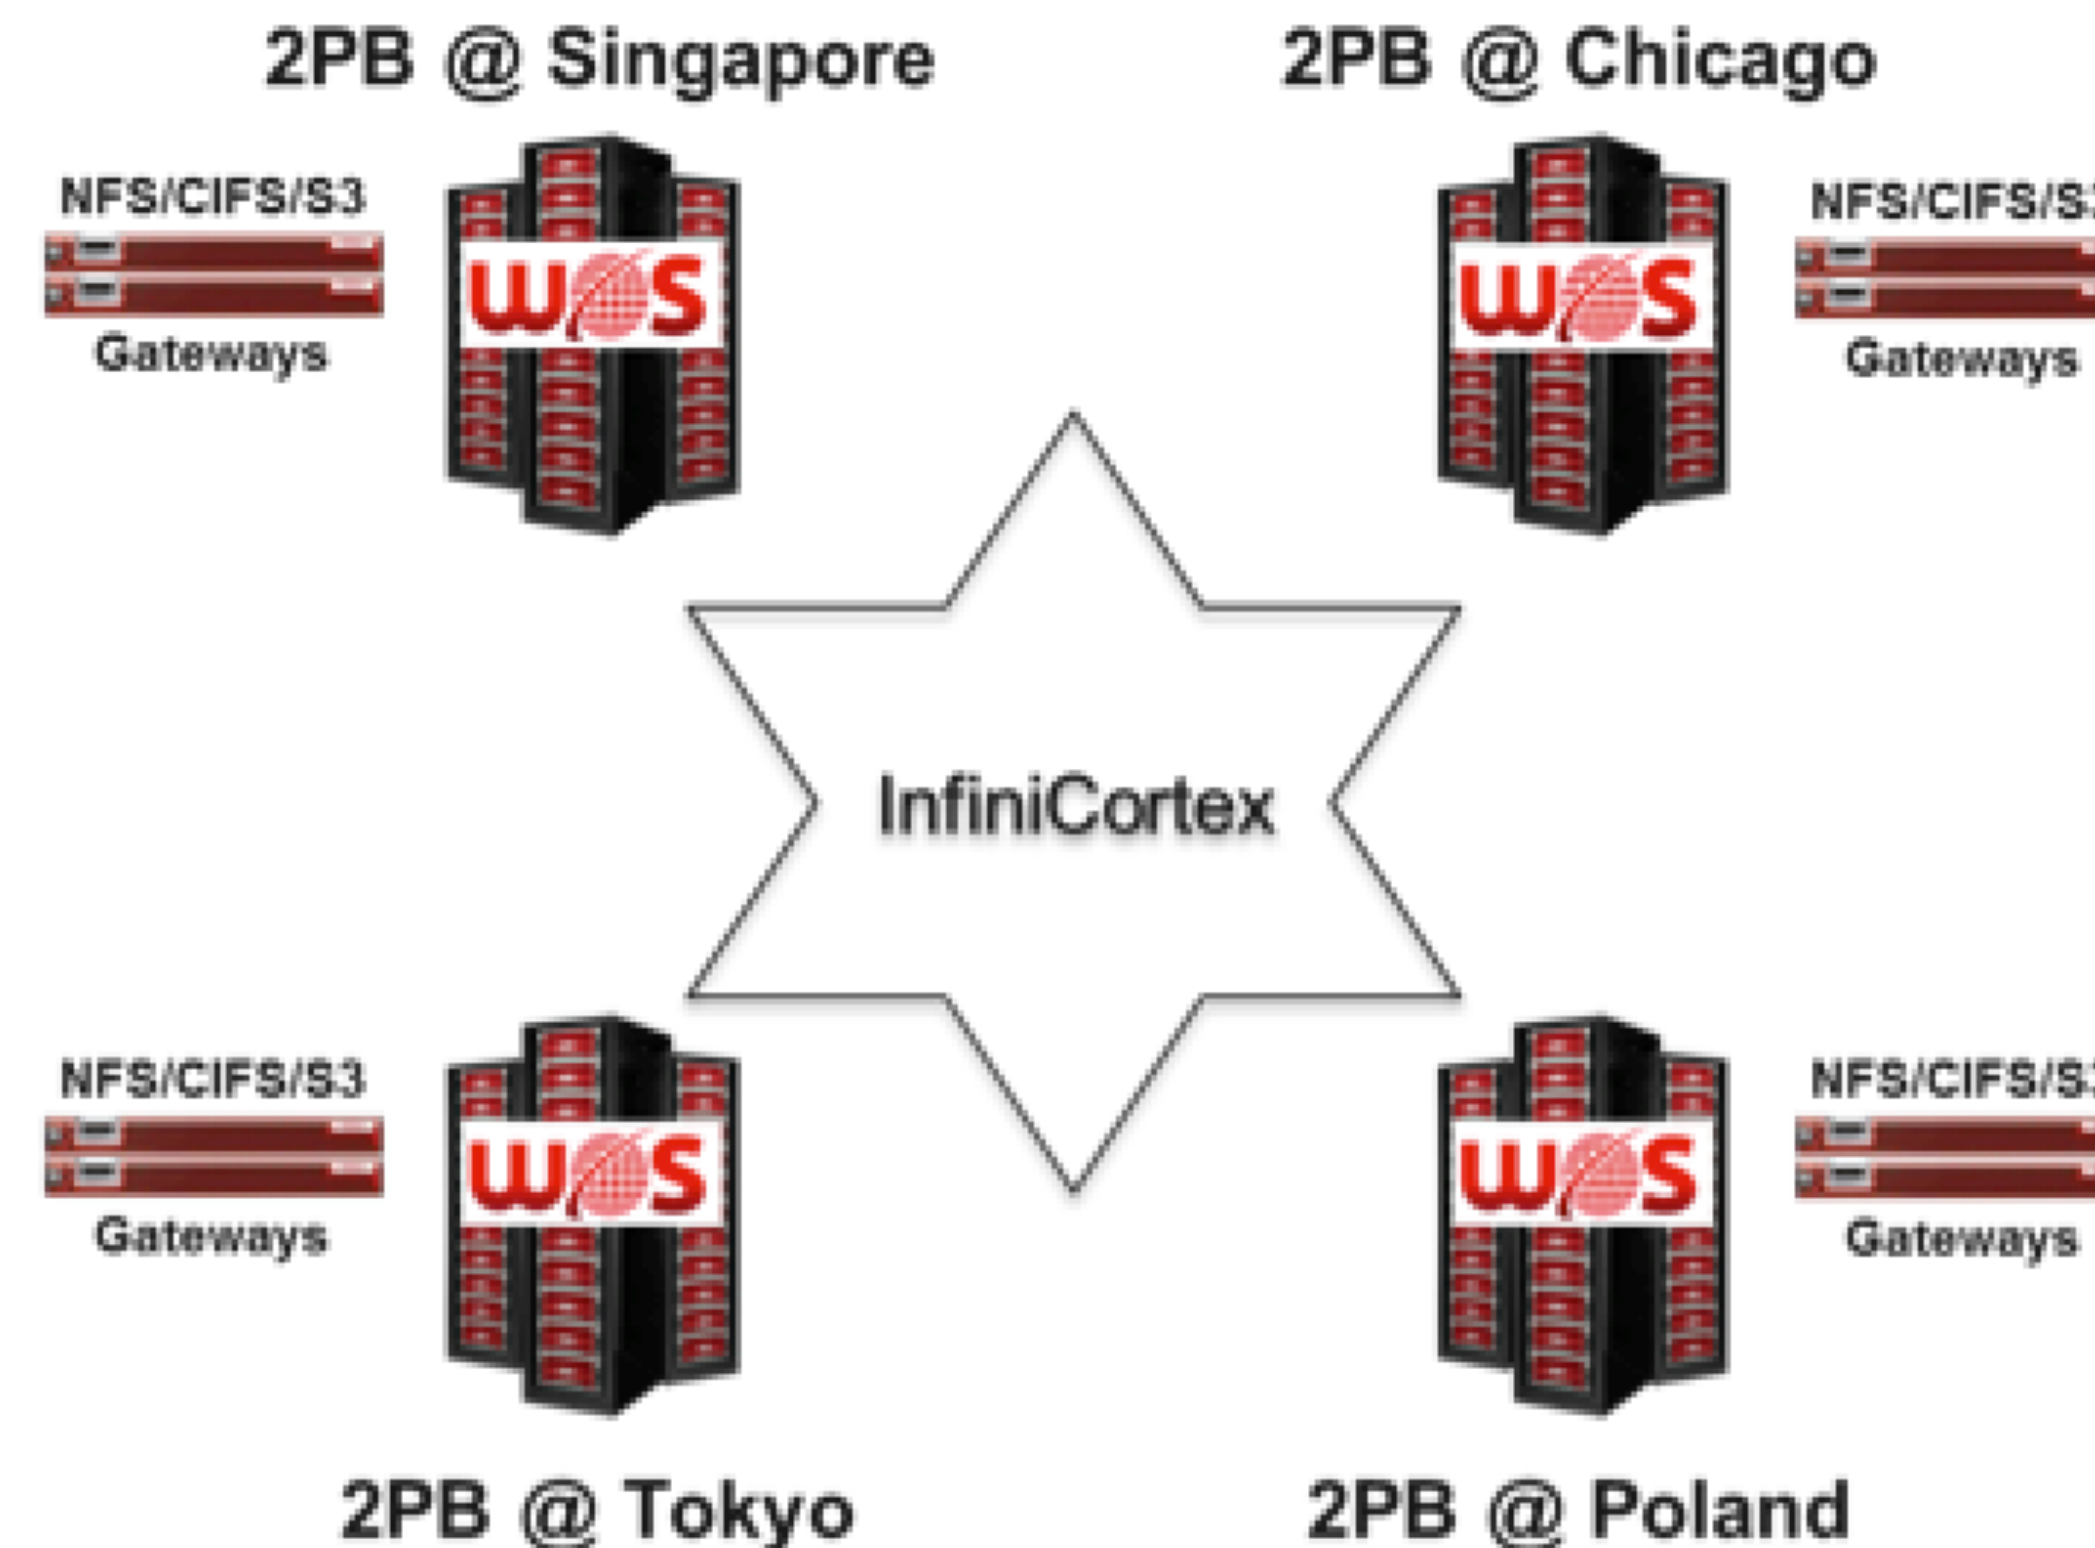
Data staging for InfiniCortex proposal

Need for InfiniCortex

DDN WOS Object storage depends on underlying high speed, low latency network for data replication between geographically distributed sites. A dedicated InfiniCortex network enables WOS to:

- Ingest data at high speed without throttling the data source.
- Replicate contents over large distance with minimal time lag.
- Researches at multiple sites to read and write same dataset without conflicts
- Distribute critical data over multiple sites for protection
- Use latest Infiniband technology natively available on WOS at long distance

Architecture



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

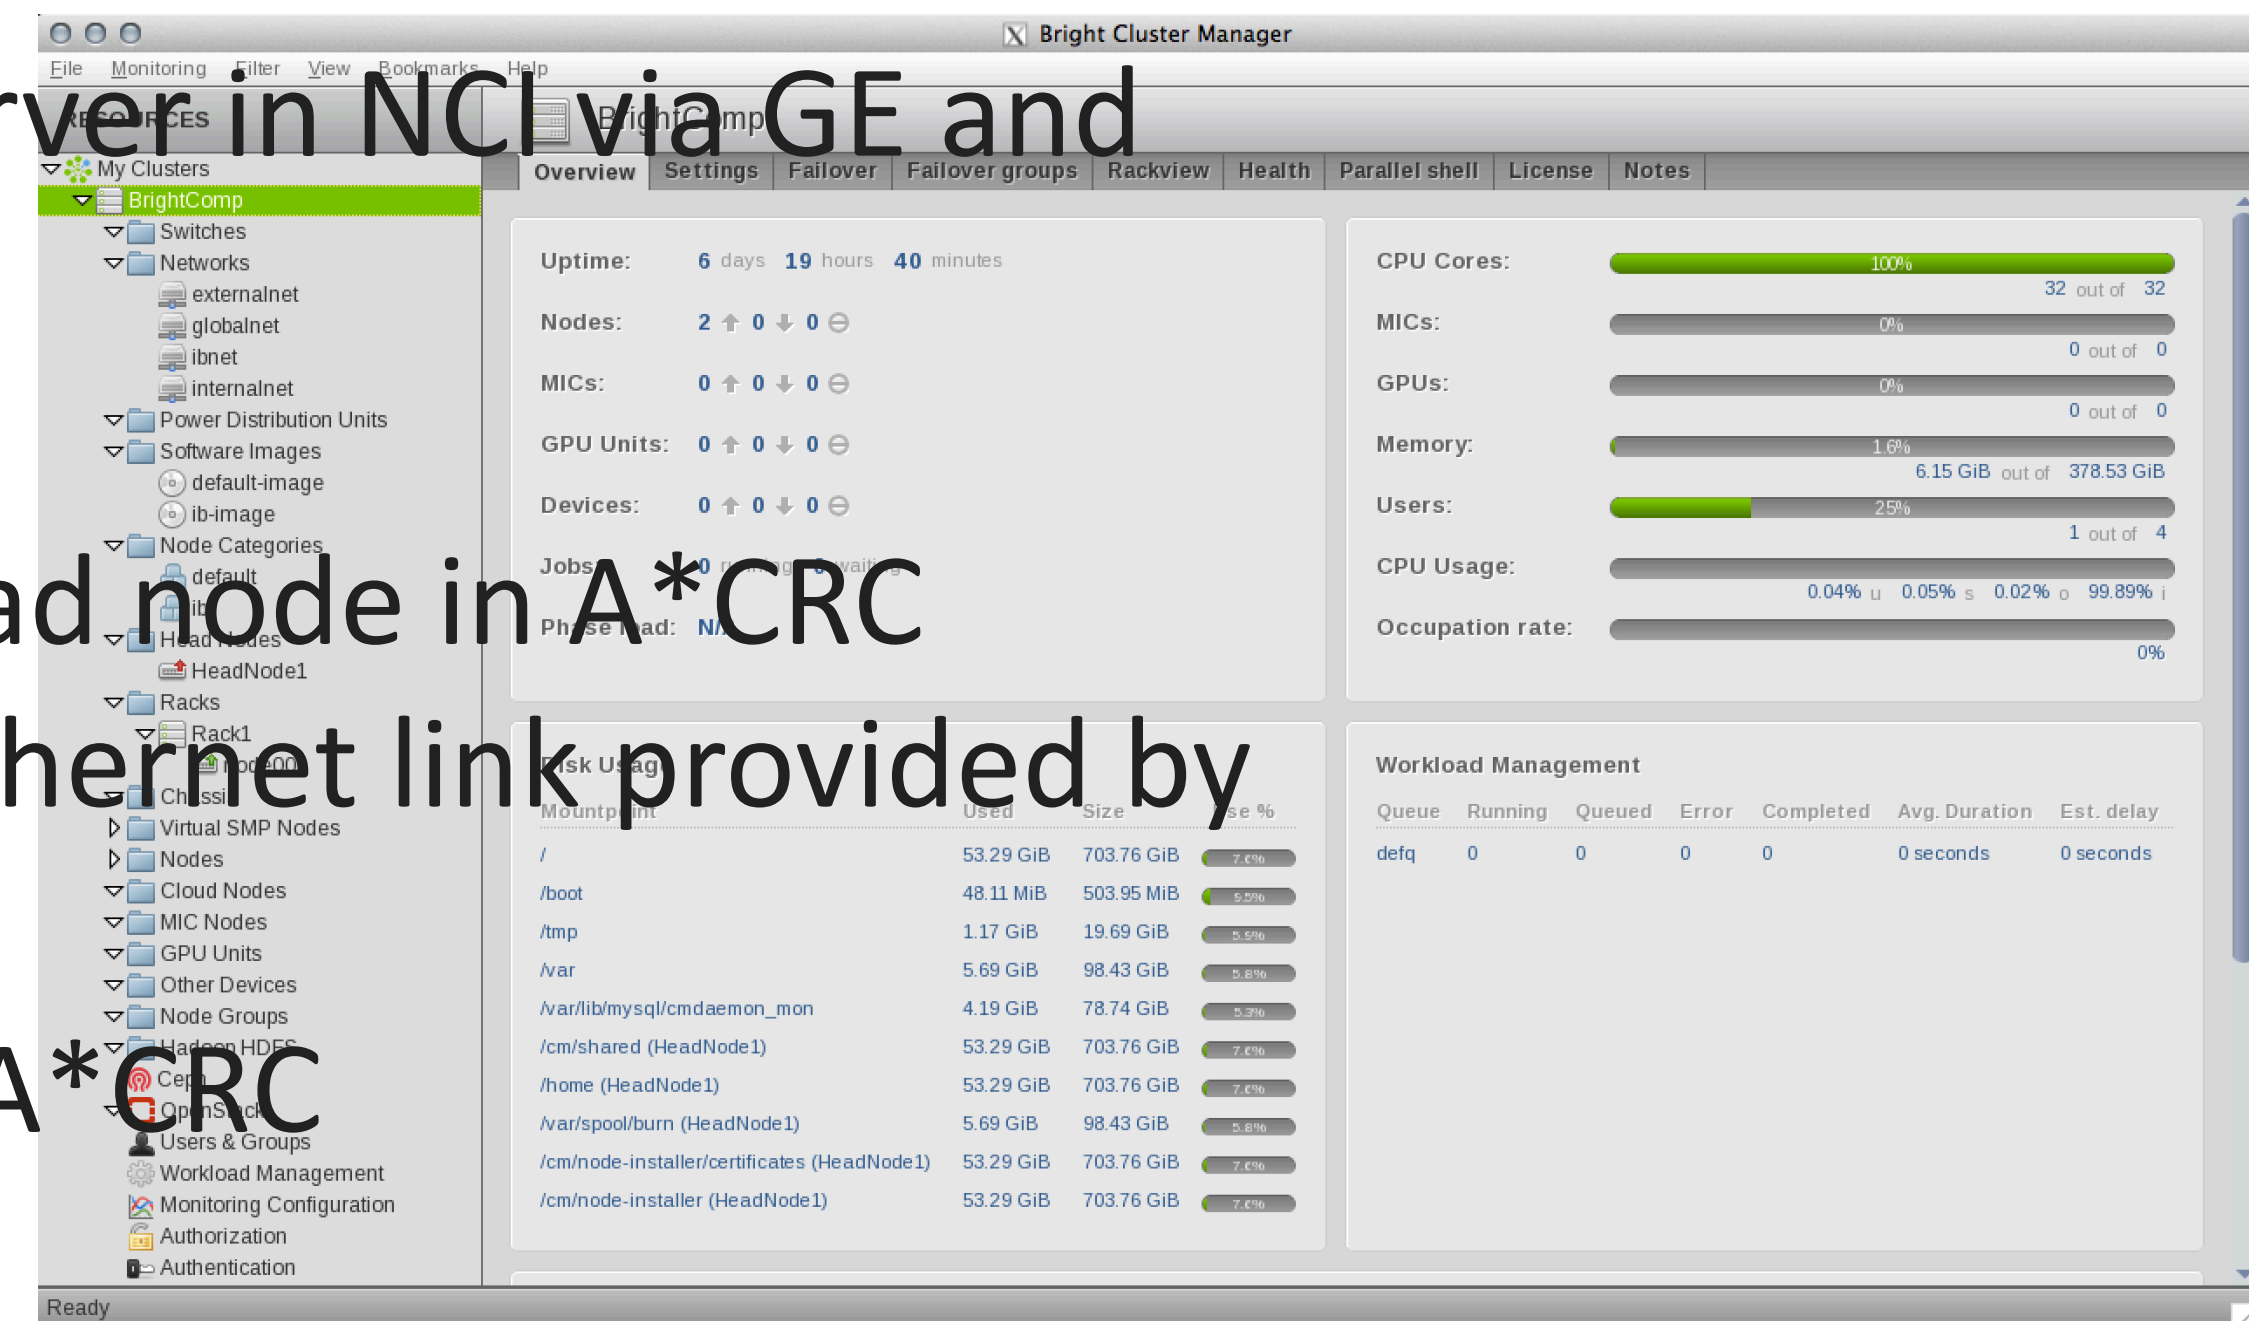
- head node in Singapore remotely provisions a server in NCI via GE and InfiniBand

Progress

- a remote node in NCI was PXE booted by the head node in A*CRC
- 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



Julia

Julia for HPC

Julia for InfiniCloud

Julia for GPUs

Julia for unums

with Prof Alan Edelman, MIT and Julia Computing, Bangalore, India

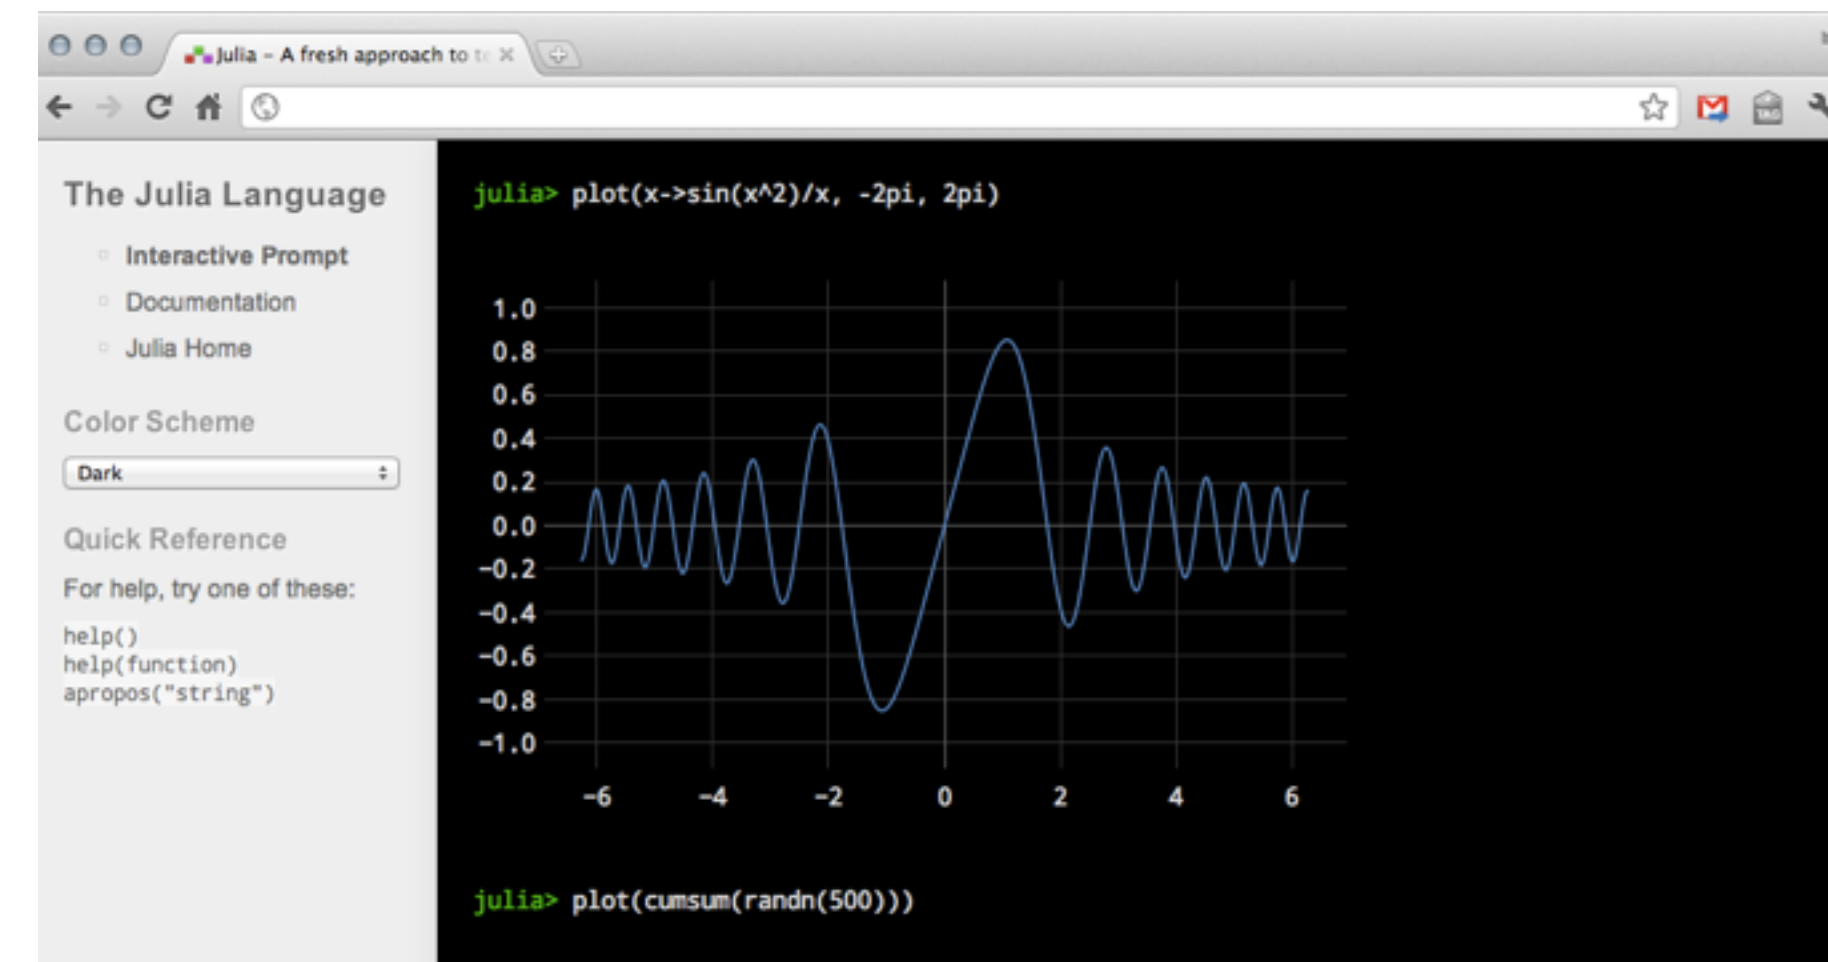
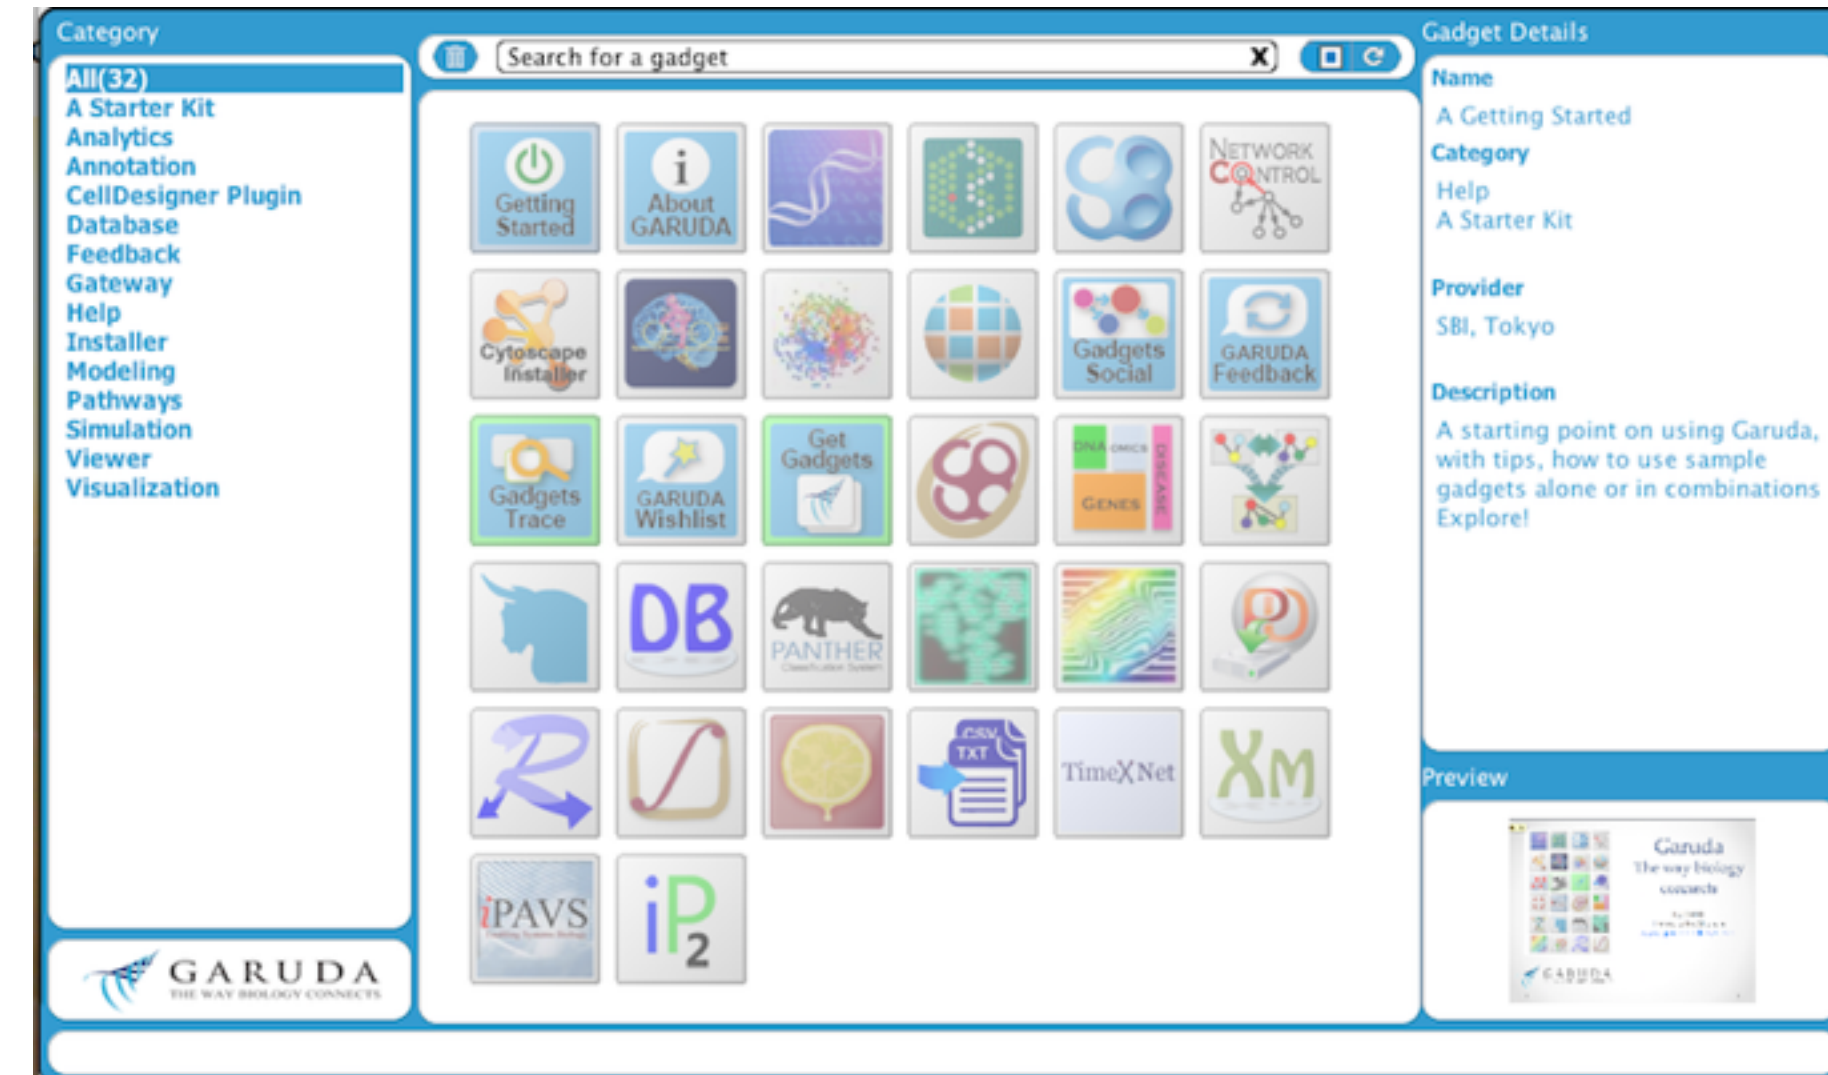
Garuda

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



unums

work with John Gustafson

starting in August 2015, initially one year

unums in Julia

unums prototyping in FGPA

Micron Automata SE Asia Exploratory Centre

work with Micron (discussions since Nov 2014)

starting in October 2015

creation of user group

tutorials and hands-on session SDK

AP hardware - 1st Q 2016

Past experience 2012

DIY Supercomputer – Technology Choices

- **CPU – Loongson**

- Further development of MIPS64
- Loongson 3 Series - R12000 with vector extensions, die shrink and multi-core packaging.
- 3A, quad-core, 800MHz, no vector, 65nm, 15W
- 3B, 8-core, 1GHz, vector, 65nm, 40W, 128GFLOPS per socket
- 3B-1500, die shrink of 3B, 1.5GHz, 256MB shared L3 cache, 32nm, 40W, faster IO interfaces, 192GFLOPS per socket
- 3C-1500, two 3B-1500 dies in one package
- 3C-2000, new layout
- 3D, HPC version, >3TFLOPS per socket

- **Network – Extoll**

- 3D torus from Heidelberg University
- Much better performance than IB FDR, and
- Multiple links

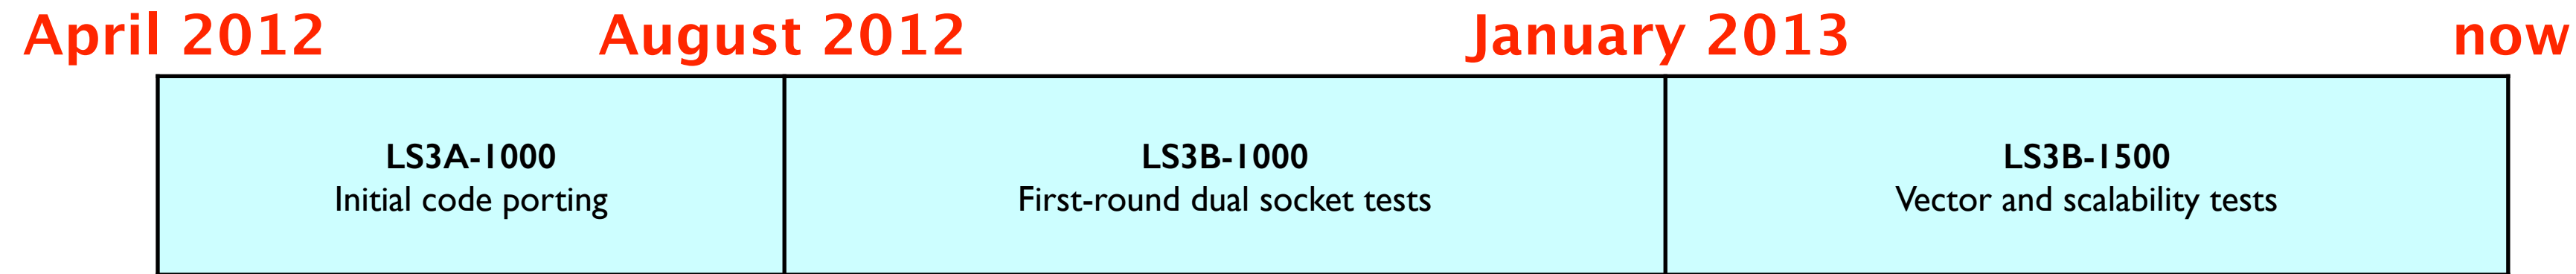


CONFIDENTIAL

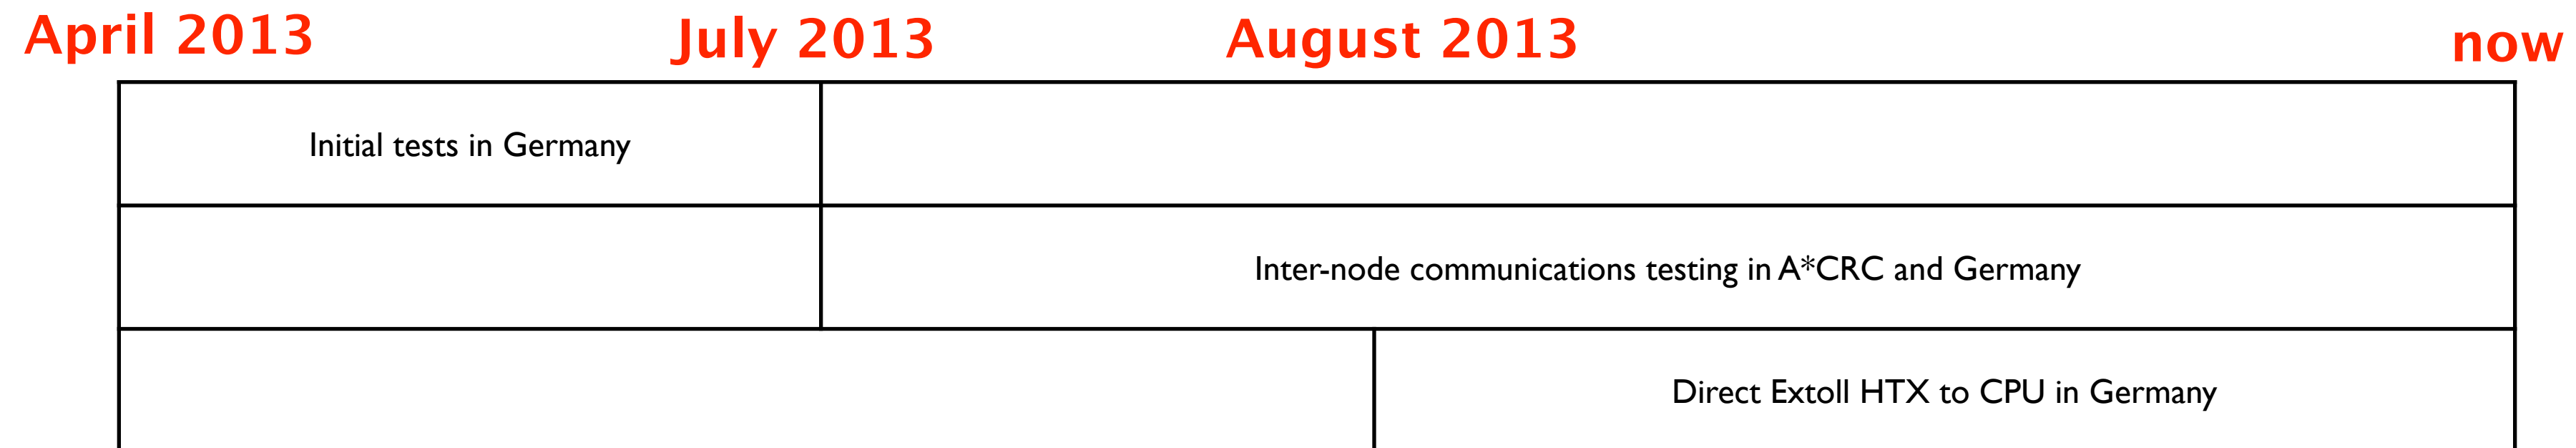
Computational
Resource Centre

Completed to Date – Hardware

- Three generations of CPU test platforms:



- Three iterations of Extoll tests on Loongson platform:



- Preparing 8-node prototype Loongson-Extoll cluster:
 - ▶ ~1.2TFLOPS
 - ▶ Will be testing for scalability and software issues
 - ▶ Whether project proceeds full scale or not will depend on the outcome of these tests

DIY Supercomputer - Stage 0 Results

- **Success**

- BLAS, CBLAS, LAPACK, LAPACKC, CORE_BLAS
- Ganglia
- Gromacs (but very slow cf. x86 – hand-coded SSE)
- Quantum Espresso
- NCBI Blast
- Octave
- PETSc
- PLASMA
- R
- Torque

- **Partial Success**

- ATLAS – failed to compile, installed via aptget
- GotoBLAS – compiled successfully, but failed tests

- **No Success**

- LAMMPS – hard-coded “-msse”
- NAMD – current version has hard-coded “-msse”, MIPS binaries of old versions available
- NWChem – needs Global Arrays



CONFIDENTIAL

Computational
Resource Centre


DIY Supercomputer – International Meetings

- **Loongson (Beijing)**
 - Feb 2012
 - Introductory meeting
 - Obtained 3A evaluation system
- **Loongson (Beijing)**
 - May 2012
 - Checkpoint meeting
 - Managed to secure in-principal agreement for 3B evaluation system
- **Shenzhen University**
 - July 2012
 - Designers of existing Loongson systems
 - Side trip to NSCC (Shenzhen) to see the only Loongson supercomputer in the world
- **A*STAR All Participants Meeting**
 - 11-12 Oct 2012



CONFIDENTIAL

Computational
Resource Centre



CONFIDENTIAL

Computational
Resource Centre

**Where will microserver
prototypes be built?**

microservers: HEAT STAR project

work with IBM Research Labs in Zurich and B Barcelona

discussions since June 2013

grant proposal to Infocom Development Authority (Green Datacenter)

results October 2015 (SGD2M)

project starts January 2016 (+ further funds raising)

refer to:

Ronald P. Luijten: SKA, DOME & Astron Project
IBM Research Zurich

~128 servers per 1U
~1,000 servers per rack

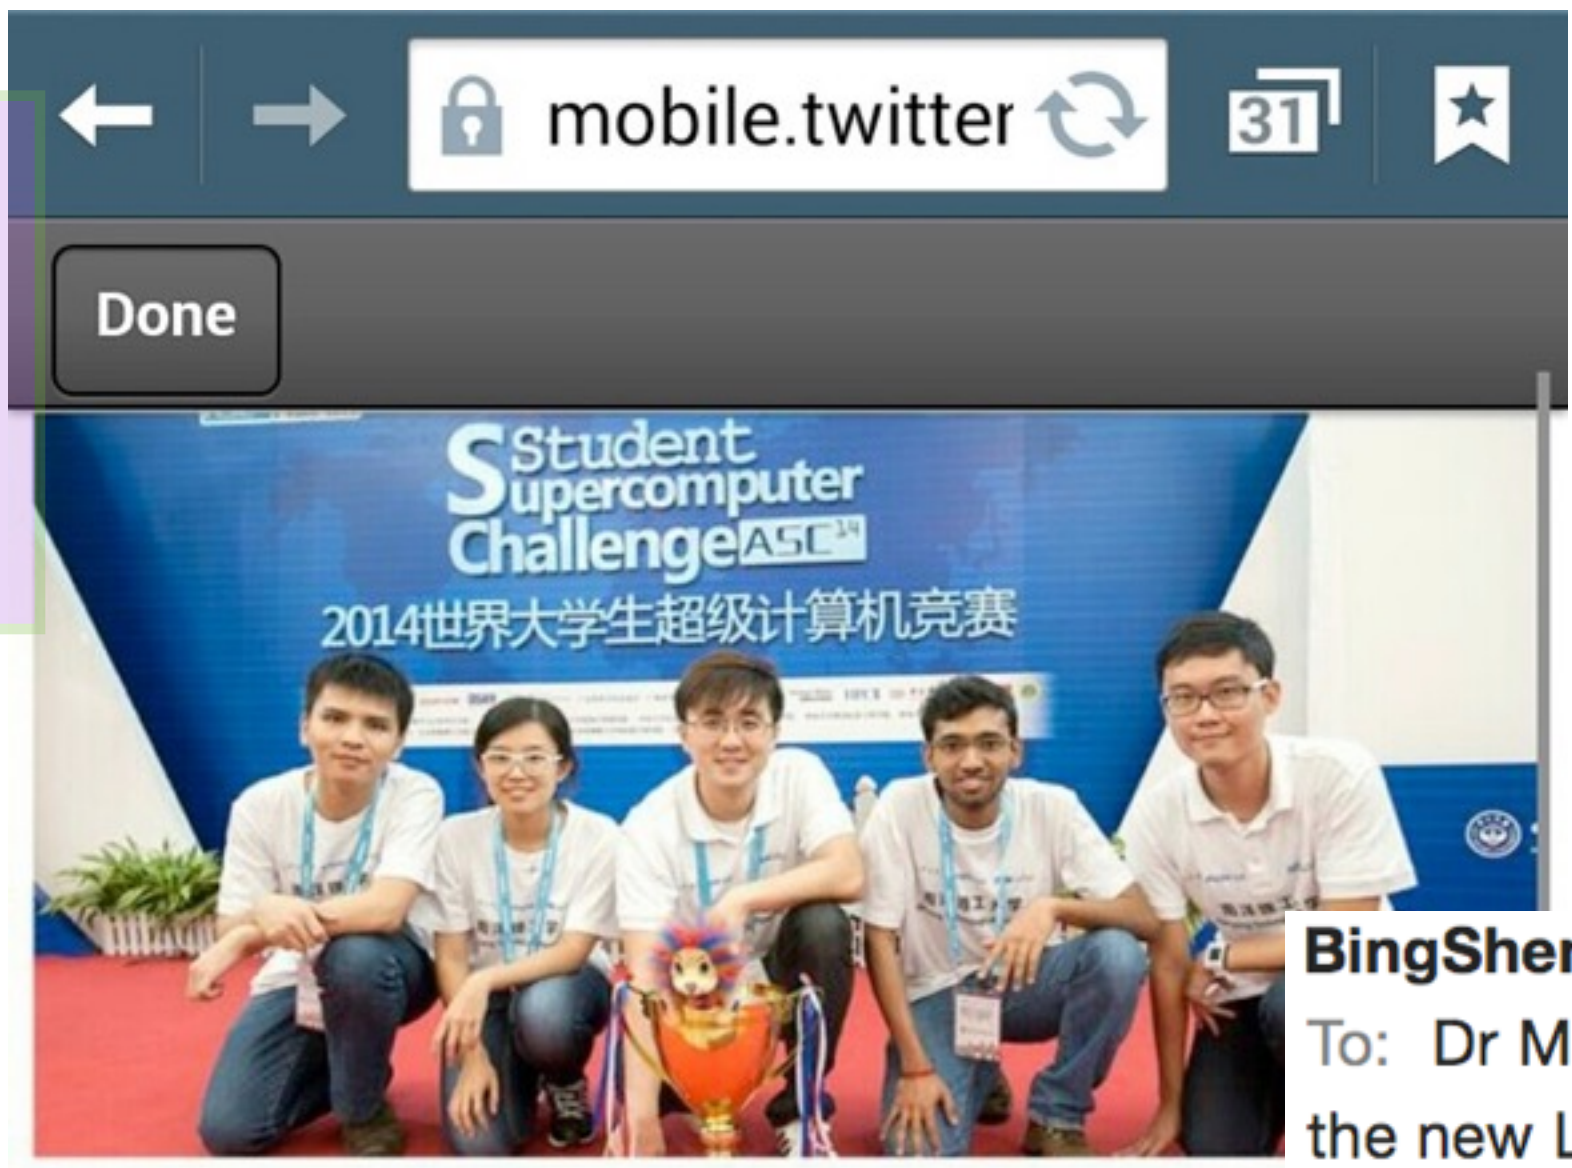
Last week, I read in Lianhe Zaobao that a team of Nanyang Technological University students had come in second out of 82 international teams, at the world's largest supercomputer competition in Guangzhou. The teams had to develop the fastest and most optimised mini-supercomputer. The participants needed many skills: computer design, engineering and coding. (The Shanghai Jiaotong University team took the first prize.) (<http://on.fb.me/1nymr9z>)

Yesterday I learnt that Shannon Lee Xinjing, an National Junior College student, had won the Intel Foundation Young Scientist Award. Shannon developed a cheaper, eco-friendly way to make rechargeable zinc-air batteries more stable - by using brinjals! The rest of our team also did well. (<http://bit.ly/1pisSPr>)

I am happy to see our students excelling in the sciences, engineering and IT. These are fascinating and exciting subjects. They will also be valuable skills to have, in the high-tech world we live and work in. I will always be glad that I studied Maths and Computer Science at university, even though I did not continue with them afterwards. I hope more students will make these subjects their passion. - LHL

[Top - NTU team in Guangzhou (L-R): Frankie Wong, Qin Gemeng, Han Jianglei, Dhakshina Moorthy Pradeep, Lim Chun Leng / Photo Courtesy of NTU]

Bottom - S'pore team in LA (L-R): Mark Lim Kit (RI); Ong Jun Yi (NUS High); Vanessa Tan Yun Han (Dunman High); Ou Yang Zhong Liang (NUS High); Shannon Lee Xinjing (NJC); Jeremy Goh Jin Quan (NUS High); Foo Cher Ying (RGS); Lim Ee Chuan Eliot (NUS High); Dousabel Tay May Yi (NUS High); Yu Shiyang (NUS High) / Photo Courtesy of MOE] - with Ong Jun Yi and 12 others.



Lee Hsien Loong @leehsienl...
Our students came in 1st & 2nd in int'l competitions recently. We - LHL bit.ly/1kcqfP5
pic.twitter.com/DF2yoD4I3C

Singapore students at competitions

BingSheng He 20 May 2015 10:01 am
To: Dr Marek Michalewicz, #HAN JIANGLEI# and 7 more...
the new Linpack record by NTU team

Hi, Marek, Prof. Deng, Kevin, Gabriel, Jonathan, Paul and Edmund, and many others in ACRC,

In Day 1 of ASC 2015, the NTU team has achieved a world record!
The news has just come out. In the HPL (Linpack) test, the NTU team has achieved 11.92 TFlops/s, which is ranked No.1 among all the teams in the ASC 2015 final. Moreover, this is a new world record. The original world record is 10.07 TFlops/s, which was created in SC'14.

Let's cheer for the team, and send our best wishes for the coming days.

Thanks for your great support to the NTU team.

thanks, saven.



Supercomputing Frontiers 2016

following on a success of

Supercomputing Frontiers 2015

(supercomputingfrontiers2015.com) we will

organise Supercomputing Frontiers 2016

second half March 2016