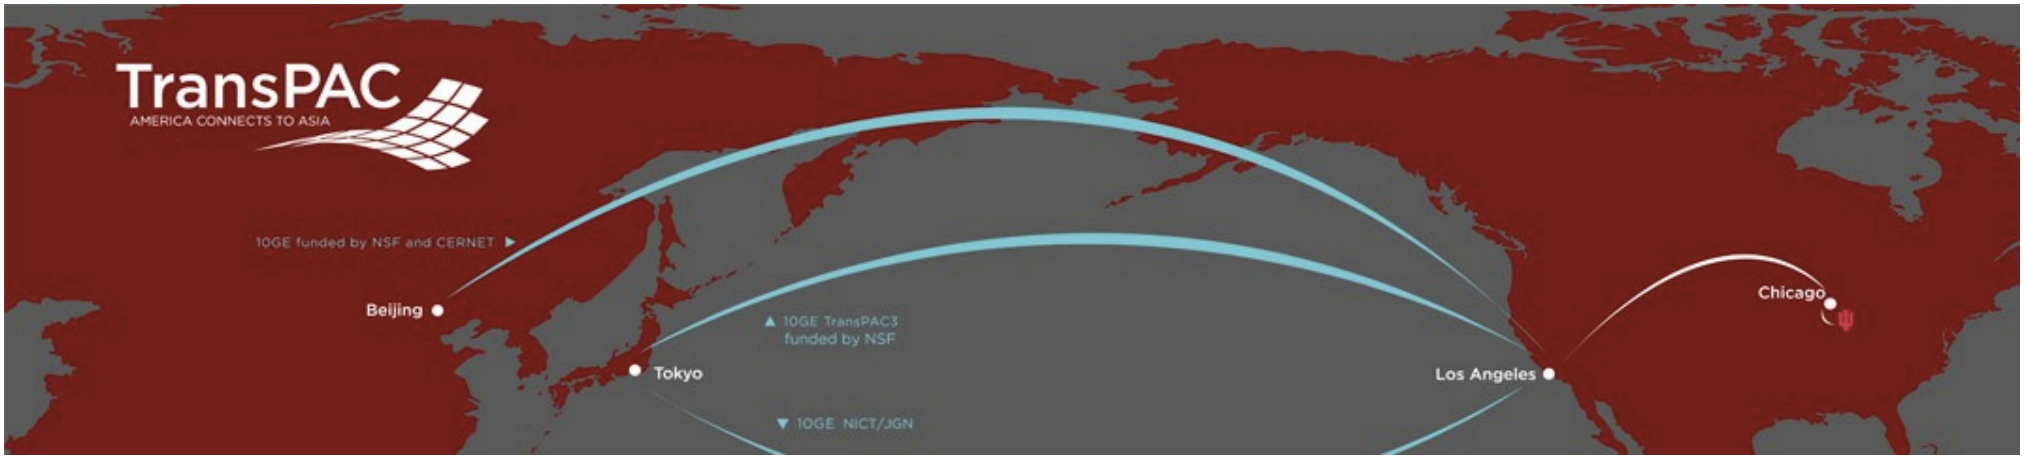

TransPAC: More than Just a Network

Andrew Lee
International Networks
Indiana University
leea@indiana.edu

Mar 17, 2016

Overview

- IN@IU
- TransPAC history
- IRNC Overview
- Network
  - Topology
  - Usage
- End Users
  - Science engagement
  - Monitoring
- Next Steps

International Networks @ Indiana University

- Group of 5
  - Dr. Jen Schopf
  - Andrew Lee
  - Hans Addleman
  - Scott Chevalier
  - Predrag Radulovic
- http://internationalnetworking.indiana.edu
- Funded primarily by the US National Science Foundation
- 3 projects:
  - TransPAC
  - ACE
  - NetSage

SC15 – Austin, TX

Plus Predrag Radulovic (Jan 4)

IRNC funded Backbones and Exchange Points

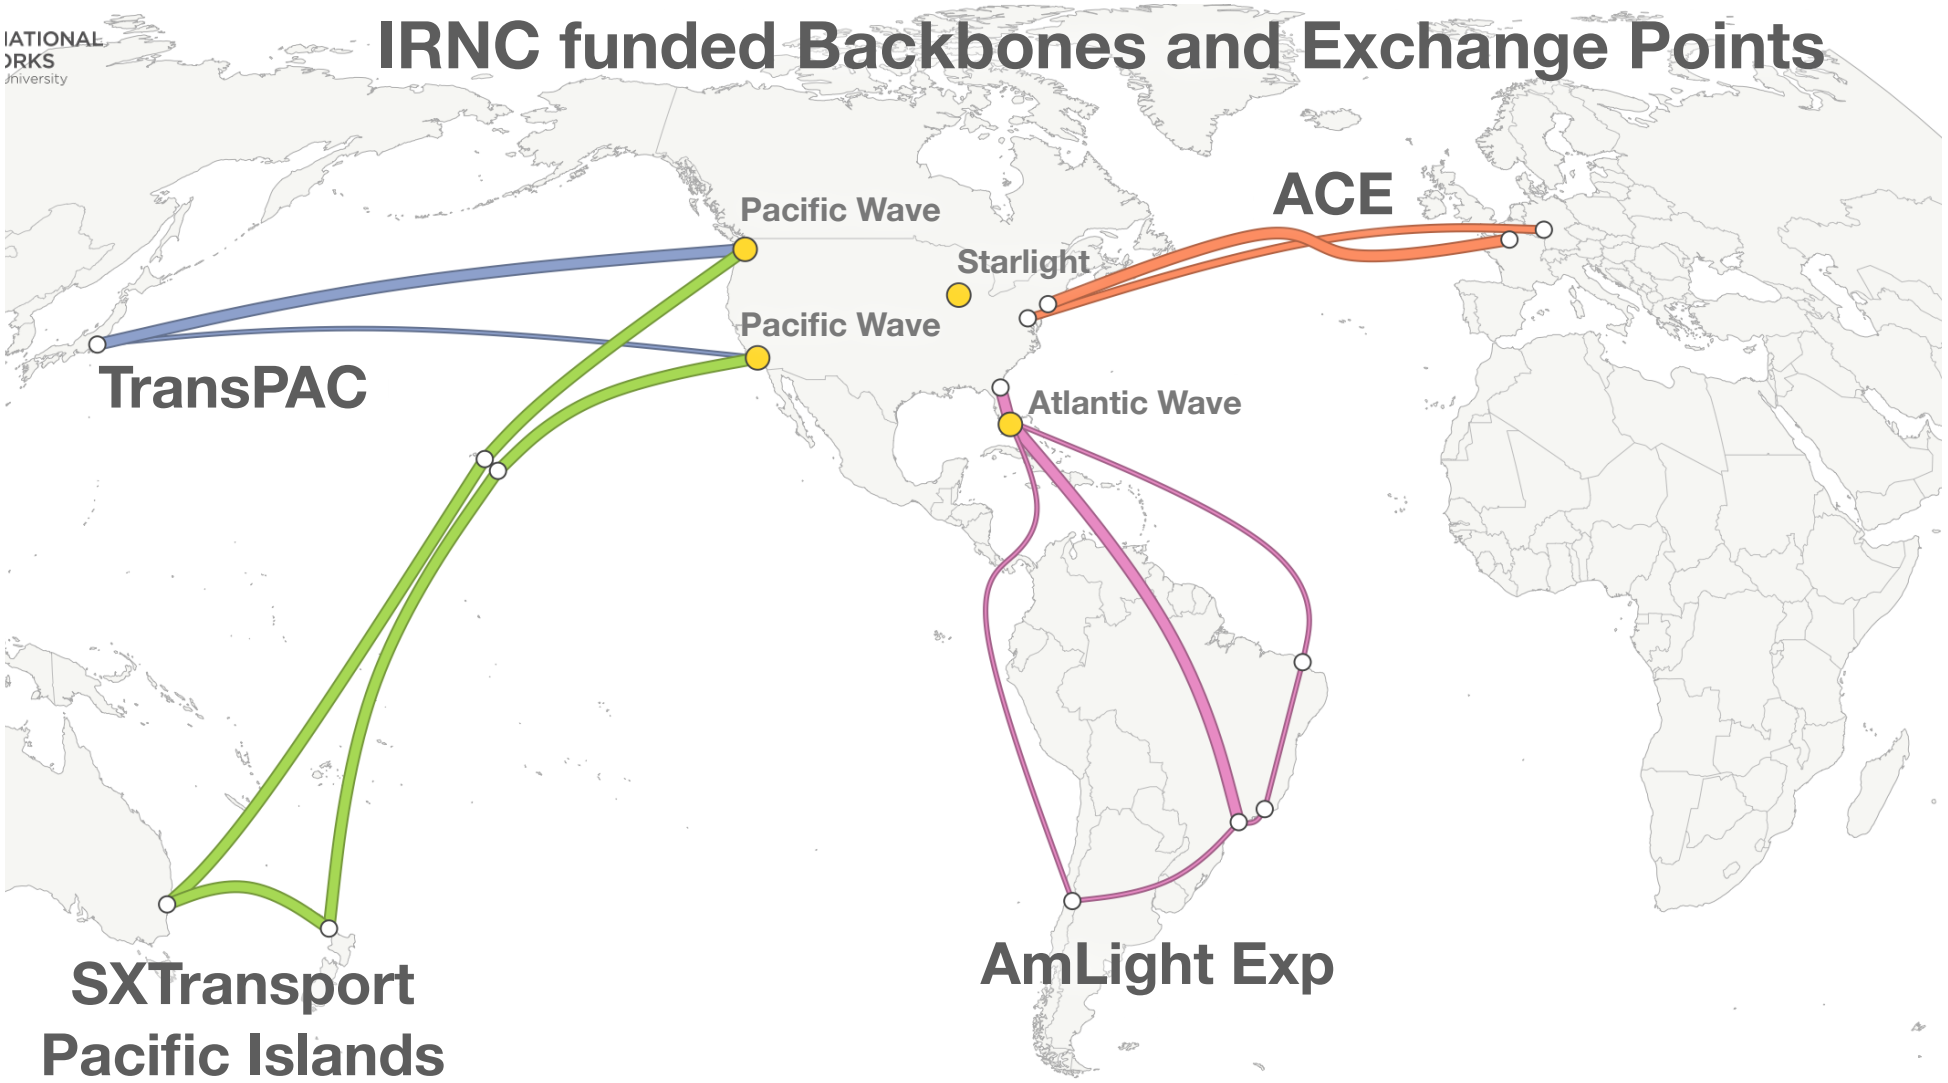

TransPAC History

- Started in 1998 by Dr. Michael McRobbie
- Cooperative partnership among Indiana University, APAN, TEIN*CC, JGN-x/NICT-Japan, NII-Japan, CERNET, and others
- Historically, 10G links to Tokyo, Beijing, others

TransPAC4 continuation

- **TransPAC4: Pragmatic Application-Driven International Networking**
- **Runs through Feb 2020**
- **$4.8M over 5 years**
- **Includes funding for circuits, exchange points, application support, research**
- **Press Release:**

https://itnews.iu.edu/articles/2015/USE-iu-globalnoc-wins-4.8-million-asia-pacific-networking-grant.php

2 10 G links LA-Tokyo

- 1 funded by NSF
  - Expires in May 2015 with end of TransPAC3
- Second funded by NICT
  - Primarily experimental use

TransPAC-Pacific Wave 100G

- Dedicated 100G circuit between Pacific Wave in Seattle and Tokyo, Japan
- Network fabric provided by Pacific Northwest Gigapop
- Points of Presence in Seattle (PacWave) and in Tokyo (WIDE/T-REX and Tata)
- 5 year agreement

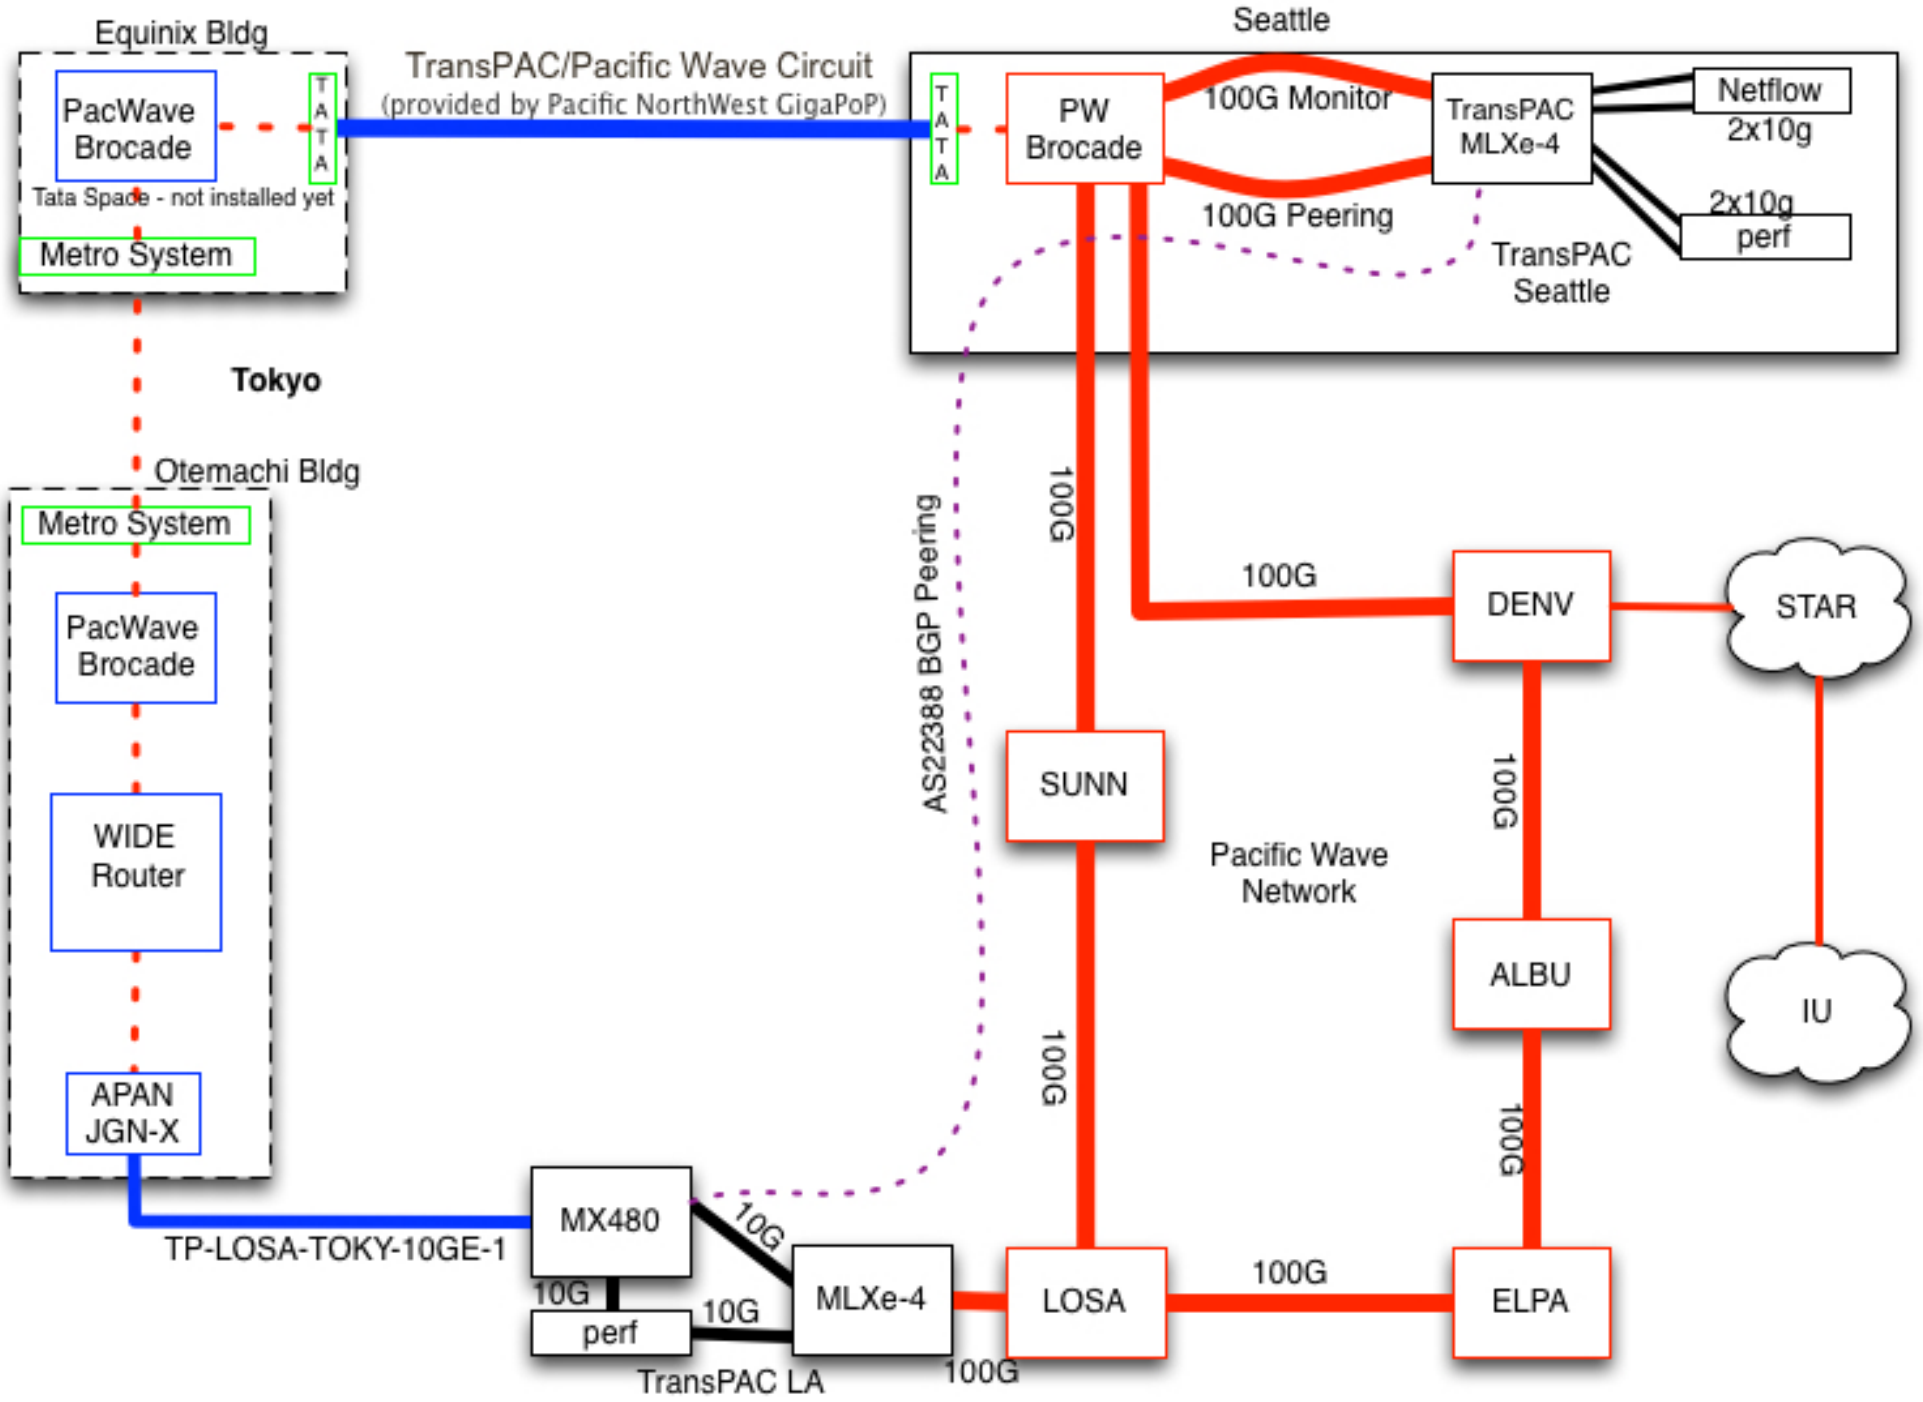
Equinix Bldg
PacWave Brocade
TATA
Tata Space - not installed yet
Metro System
TransPAC/Pacific Wave Circuit
(provided by Pacific NorthWest GigaPoP)
Seattle
TATA
PW Brocade
100G Monitor
100G Peering
TransPAC MLXe-4
Netflow 2x10g
2x10g perf
TransPAC Seattle
Tokyo
Otemachi Bldg
Metro System
PacWave Brocade
WIDE Router
APAN JGN-X
TP-LOSA-TOKY-10GE-1
AS22388 BGP Peering
100G
SUNN
DENV
STAR
Pacific Wave Network
ALBU
IU
MX480
10G
MLXe-4
10G
perf
10G
TransPAC LA
LOSA
100G
ELPA
100G

Users

- Usage Policy – all R&E traffic allowed
  - Solely commercial traffic prohibited*
  - Peering/routing to shift gradually, starting with discussion at APAN in January
- TransPAC is supporting end-user engagement to increase efficient use of the networks
- Example application from SC'15:

University of Tokyo Test over 100G ultra long distance links

- Investigating what happens to TCP performance over ultra long, ultra large circuits
  - Cubic-TCP
  - Jumbo Frames
- CPU Load
- Limitations of the TCP Protocol
  - 1GB Max buffer (Linux kernel limitation)
  - 32 bit sequence numbers (RFC 793)

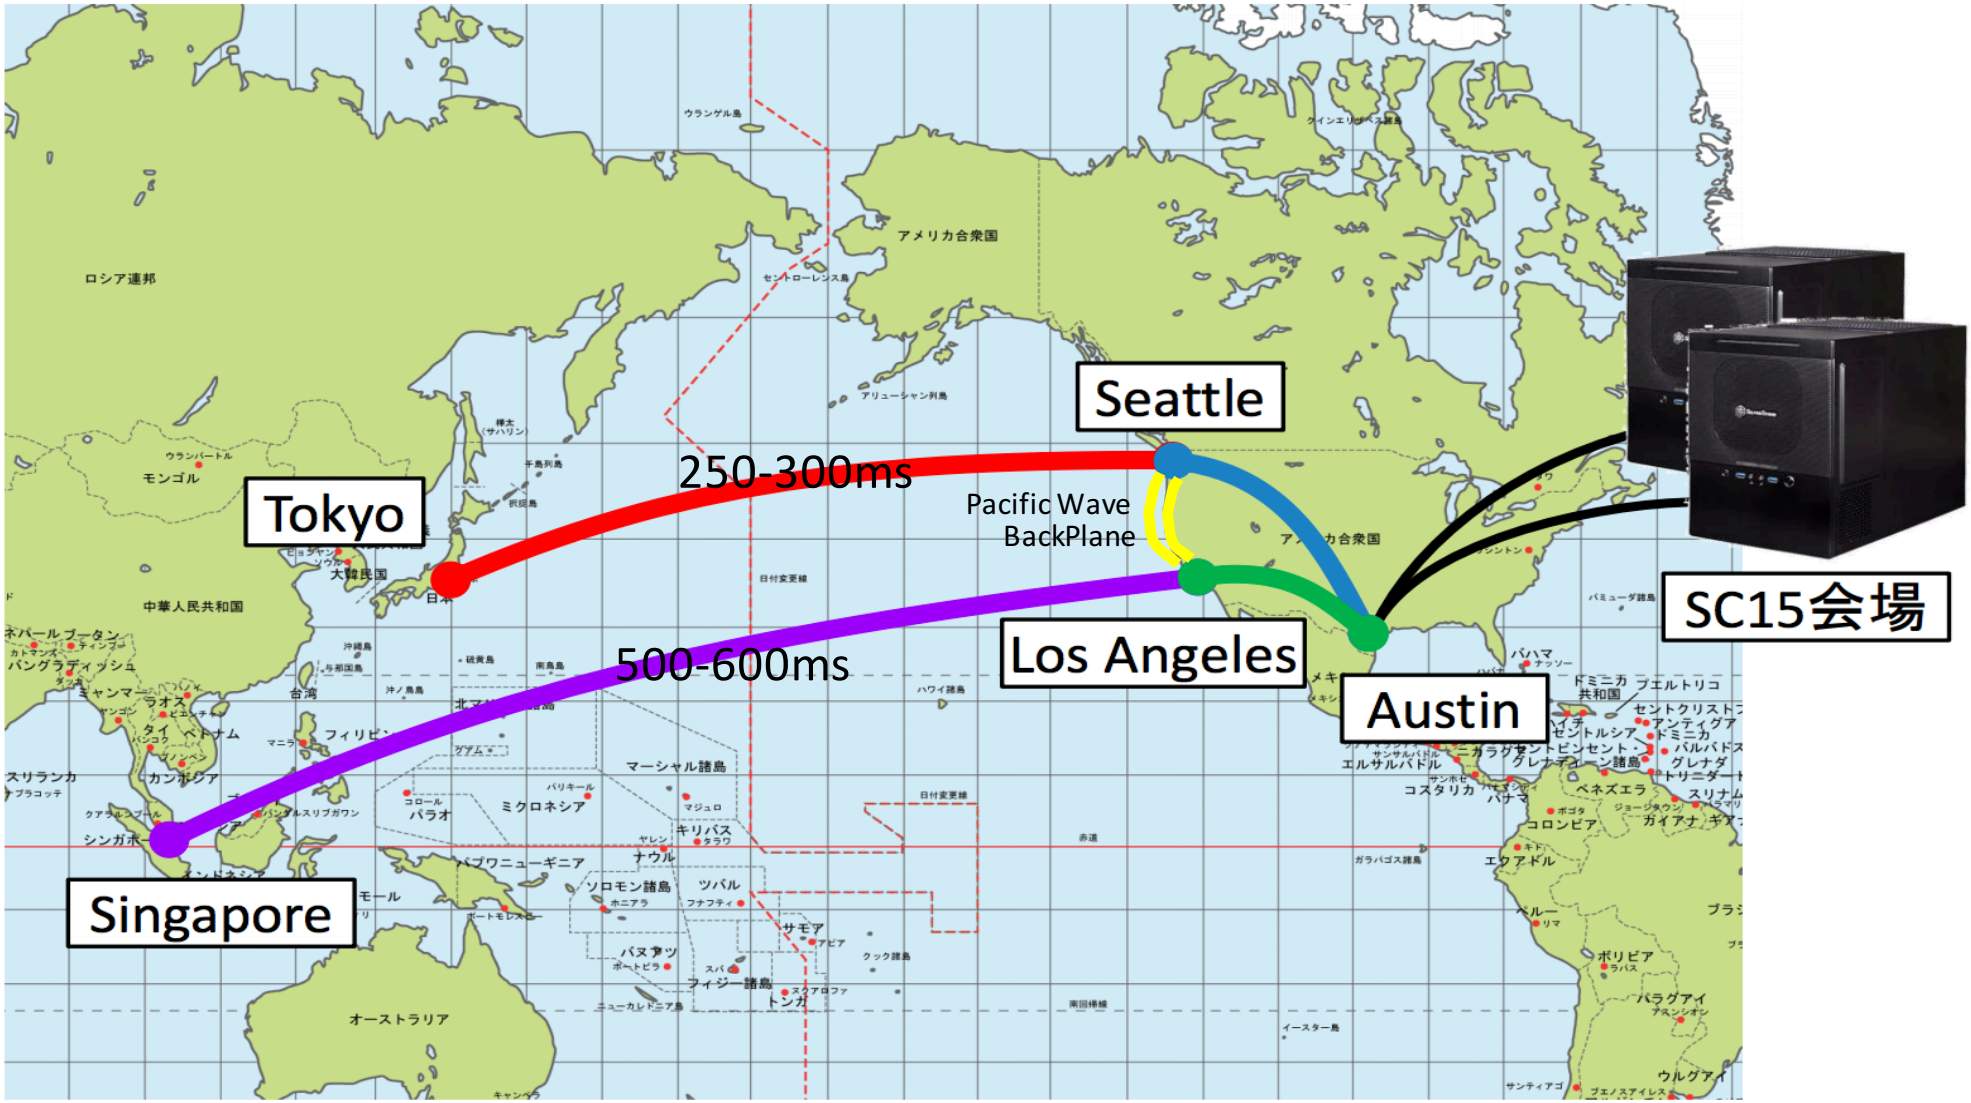

University of Tokyo TCP Experiment at SC'15

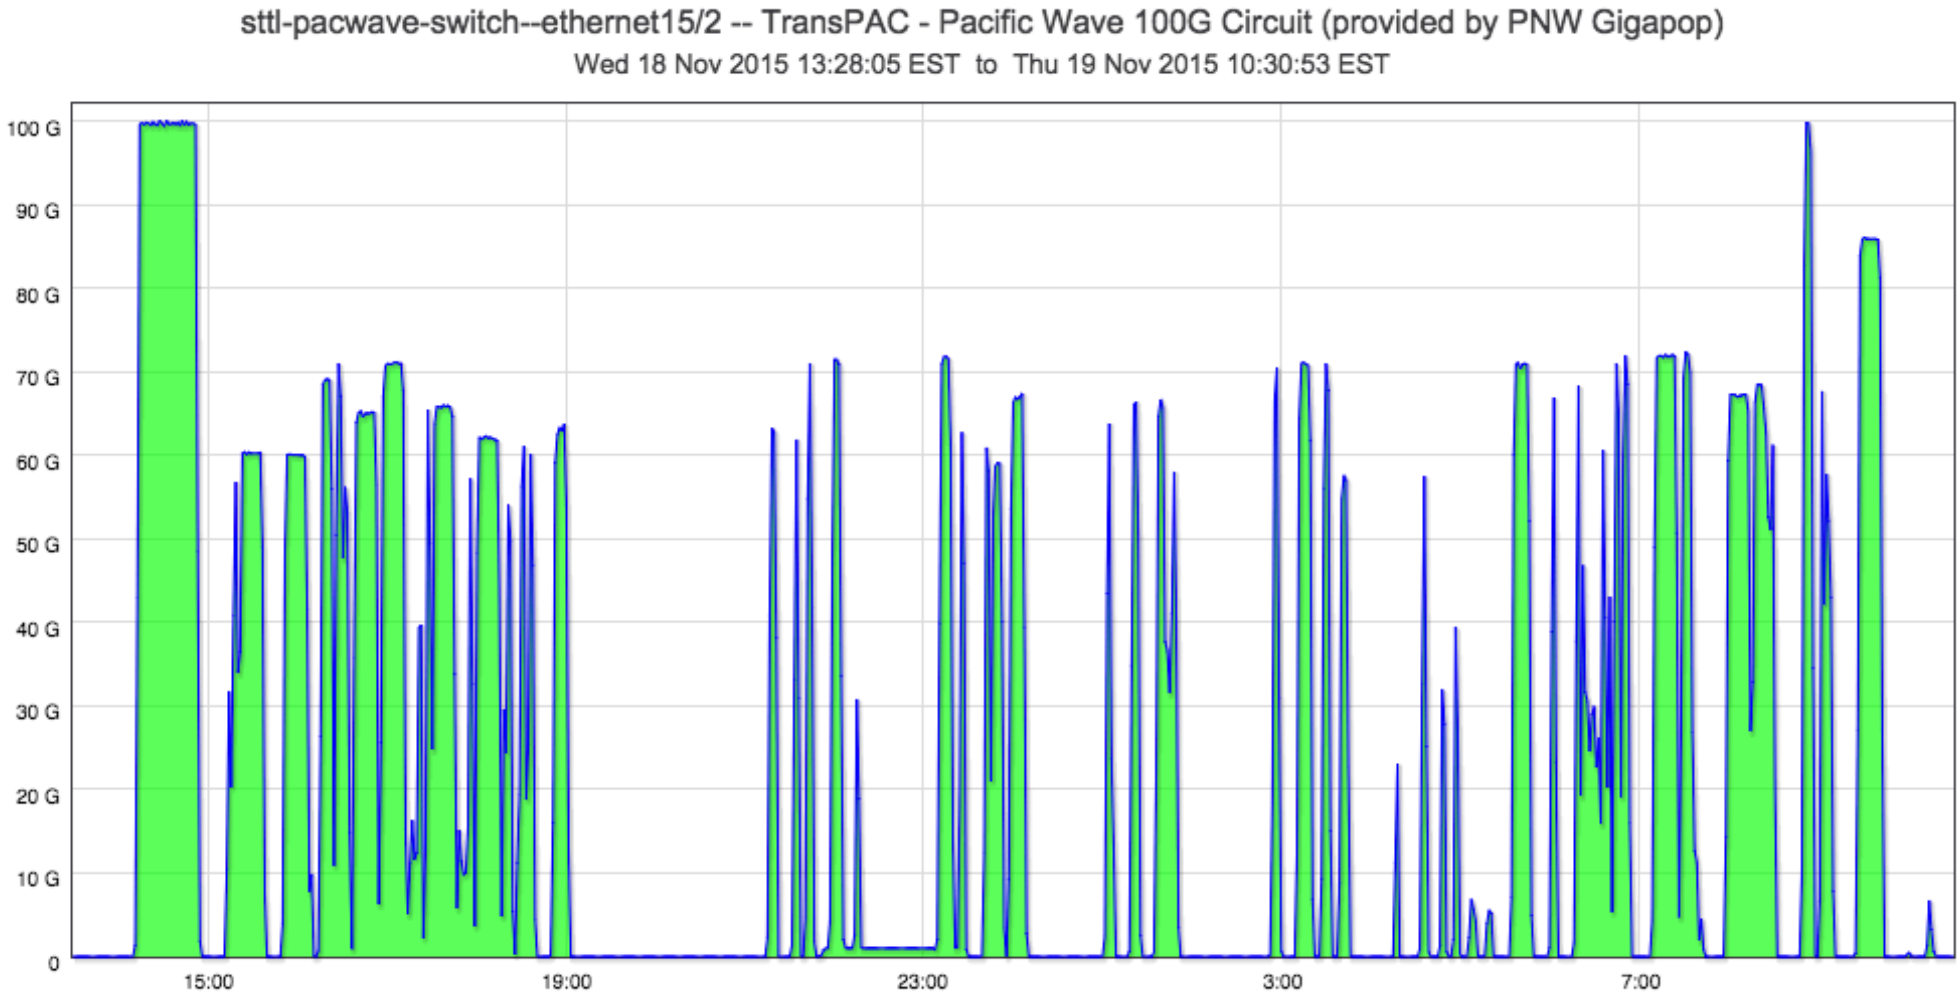
sttl-pacwave-switch--ethernet15/2 -- TransPAC - Pacific Wave 100G Circuit (provided by PNW Gigapop)
Wed 18 Nov 2015 13:28:05 EST to Thu 19 Nov 2015 10:30:53 EST
100 G
90 G
80 G
70 G
60 G
50 G
40 G
30 G
20 G
10 G
0
15:00
19:00
23:00
3:00
7:00

Who Uses TransPAC?

- De-identified flow data analysis
  - Sampled netflow v5 (no 4 byte ASNs, v6)
- 10G link only right now

Who Uses TransPAC?

- **Many of you (even if you don't realize it)**

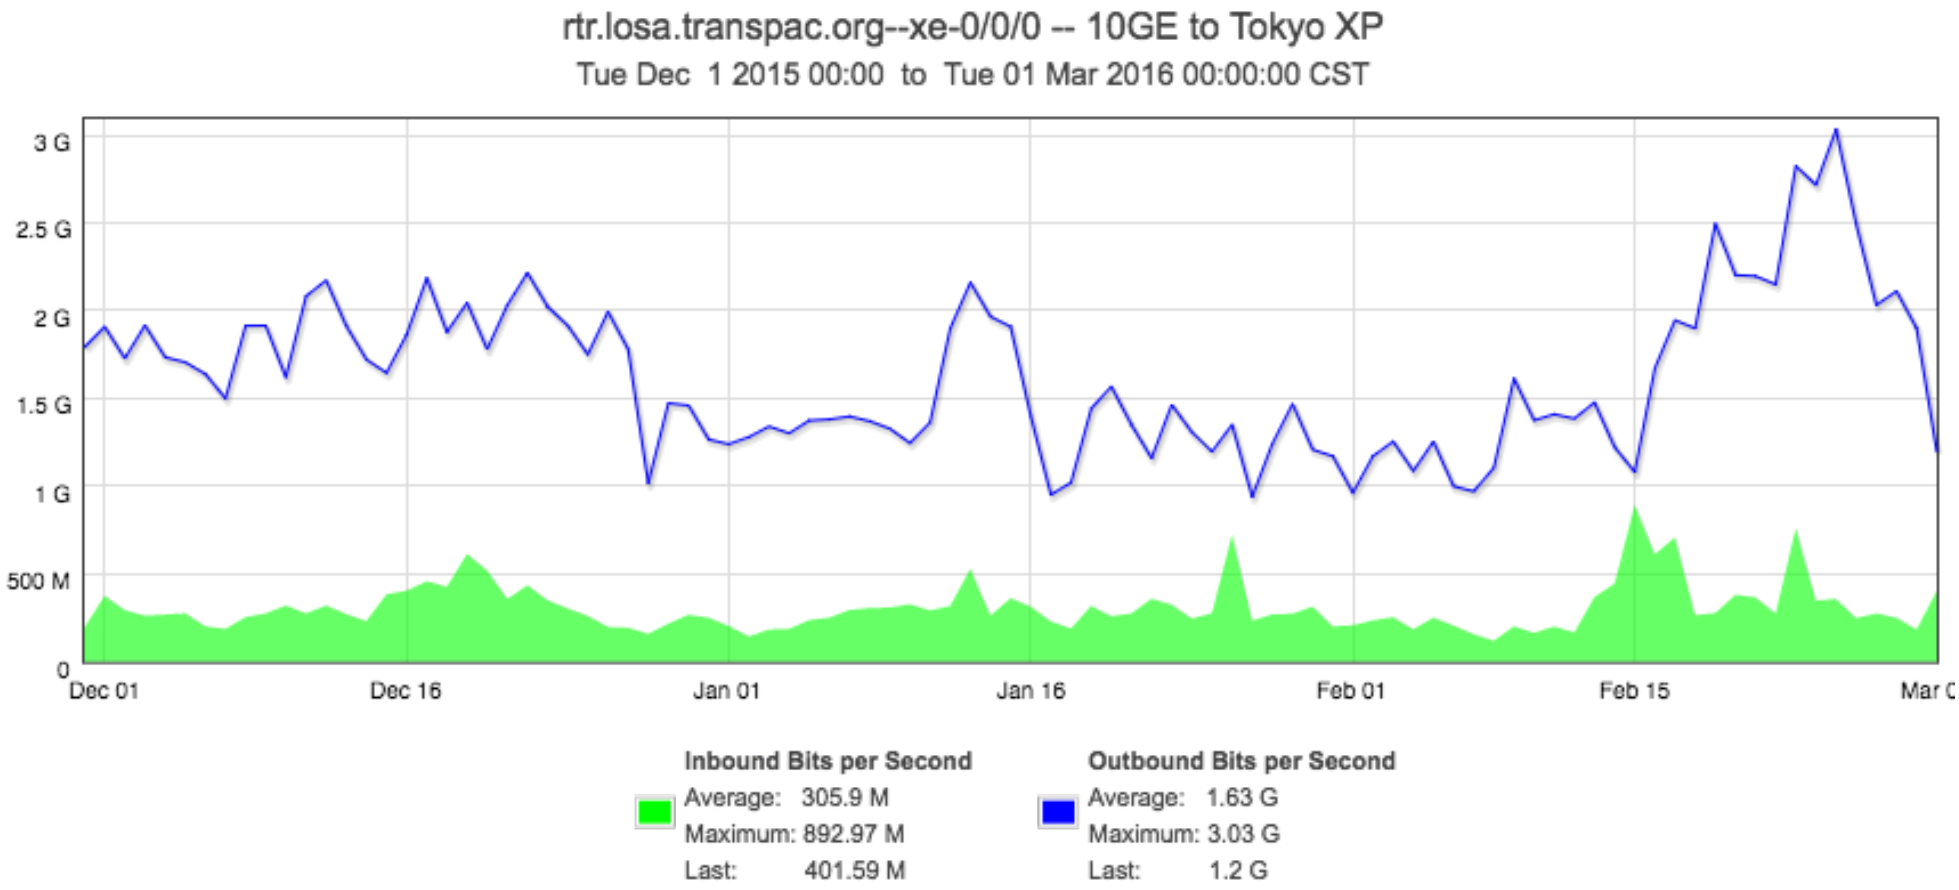
rtr.losa.transpac.org--xe-0/0/0 -- 10GE to Tokyo XP
Tue Dec 1 2015 00:00 to Tue 01 Mar 2016 00:00:00 CST
3 G
2.5 G
2 G
1.5 G
1 G
500 M
0
Dec 01
Dec 16
Jan 01
Jan 16
Feb 01
Feb 15
Inbound Bits per Second
Average: 305.9 M
Maximum: 892.97 M
Last: 401.59 M
Outbound Bits per Second
Average: 1.63 G
Maximum: 3.03 G
Last: 1.2 G

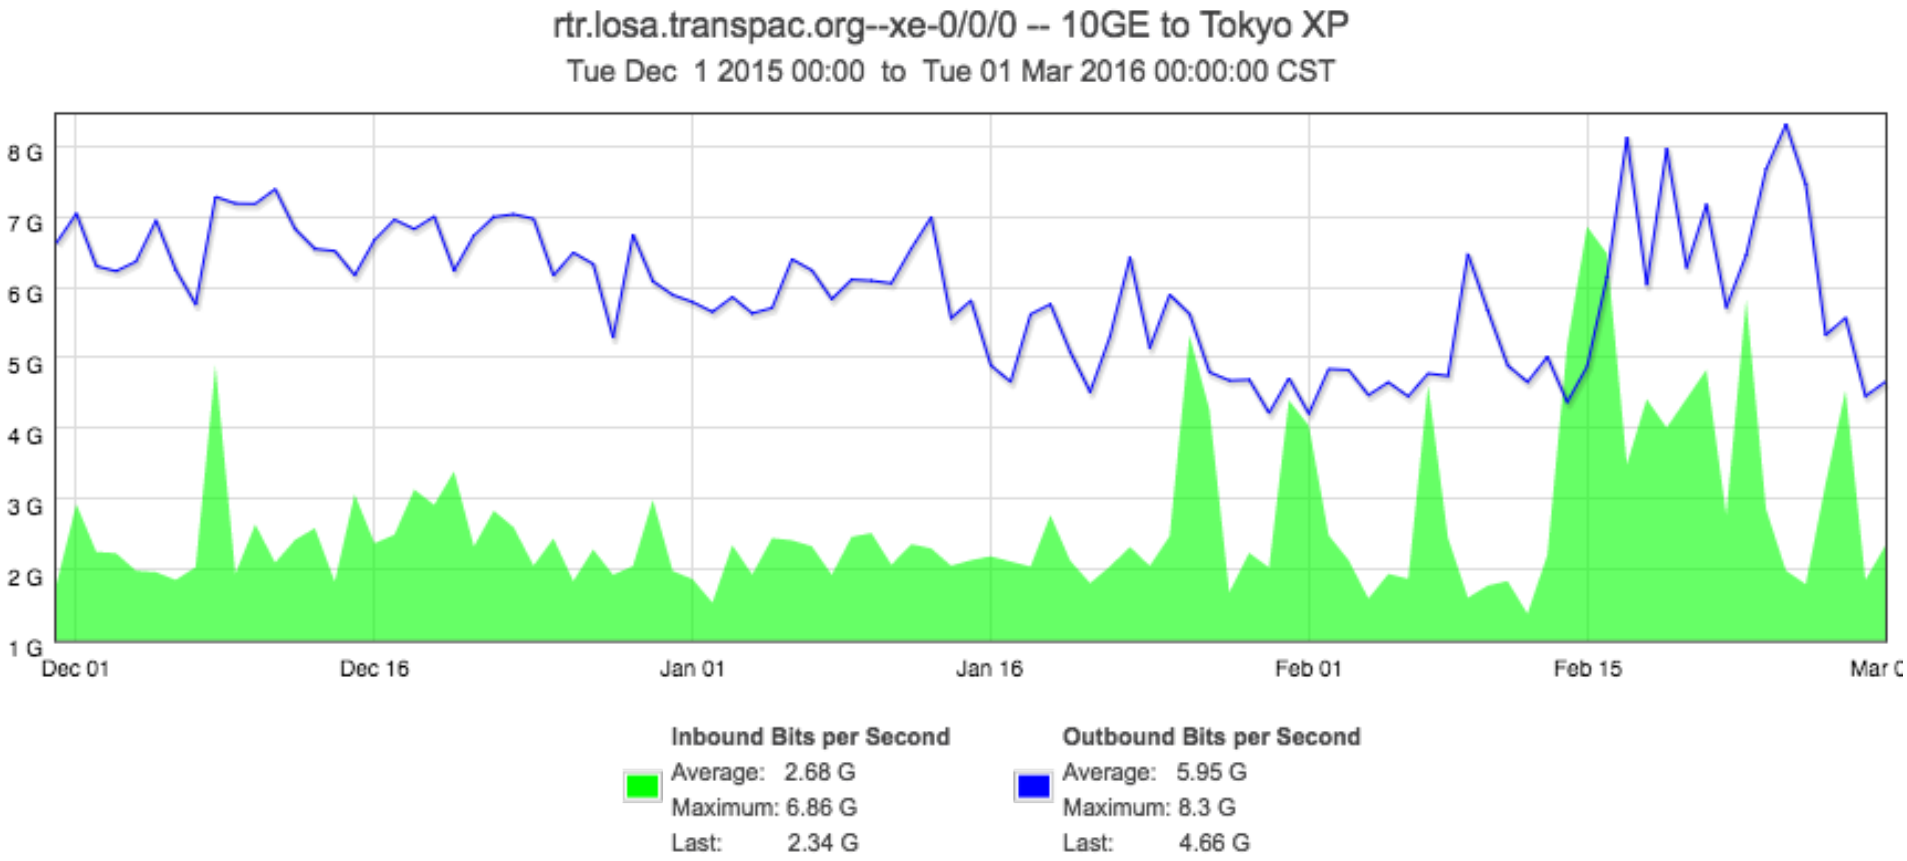
rtr.losa.transpac.org--xe-0/0/0 -- 10GE to Tokyo XP
Tue Dec 1 2015 00:00 to Tue 01 Mar 2016 00:00:00 CST
8 G
7 G
6 G
5 G
4 G
3 G
2 G
1 G
Dec 01
Dec 16
Jan 01
Jan 16
Feb 01
Feb 15
Inbound Bits per Second
Average: 2.68 G
Maximum: 6.86 G
Last: 2.34 G
Outbound Bits per Second
Average: 5.95 G
Maximum: 8.3 G
Last: 4.66 G

TransPAC4 Year 2 Quarter 1 Top Ten Talkers by Source AS Outbound

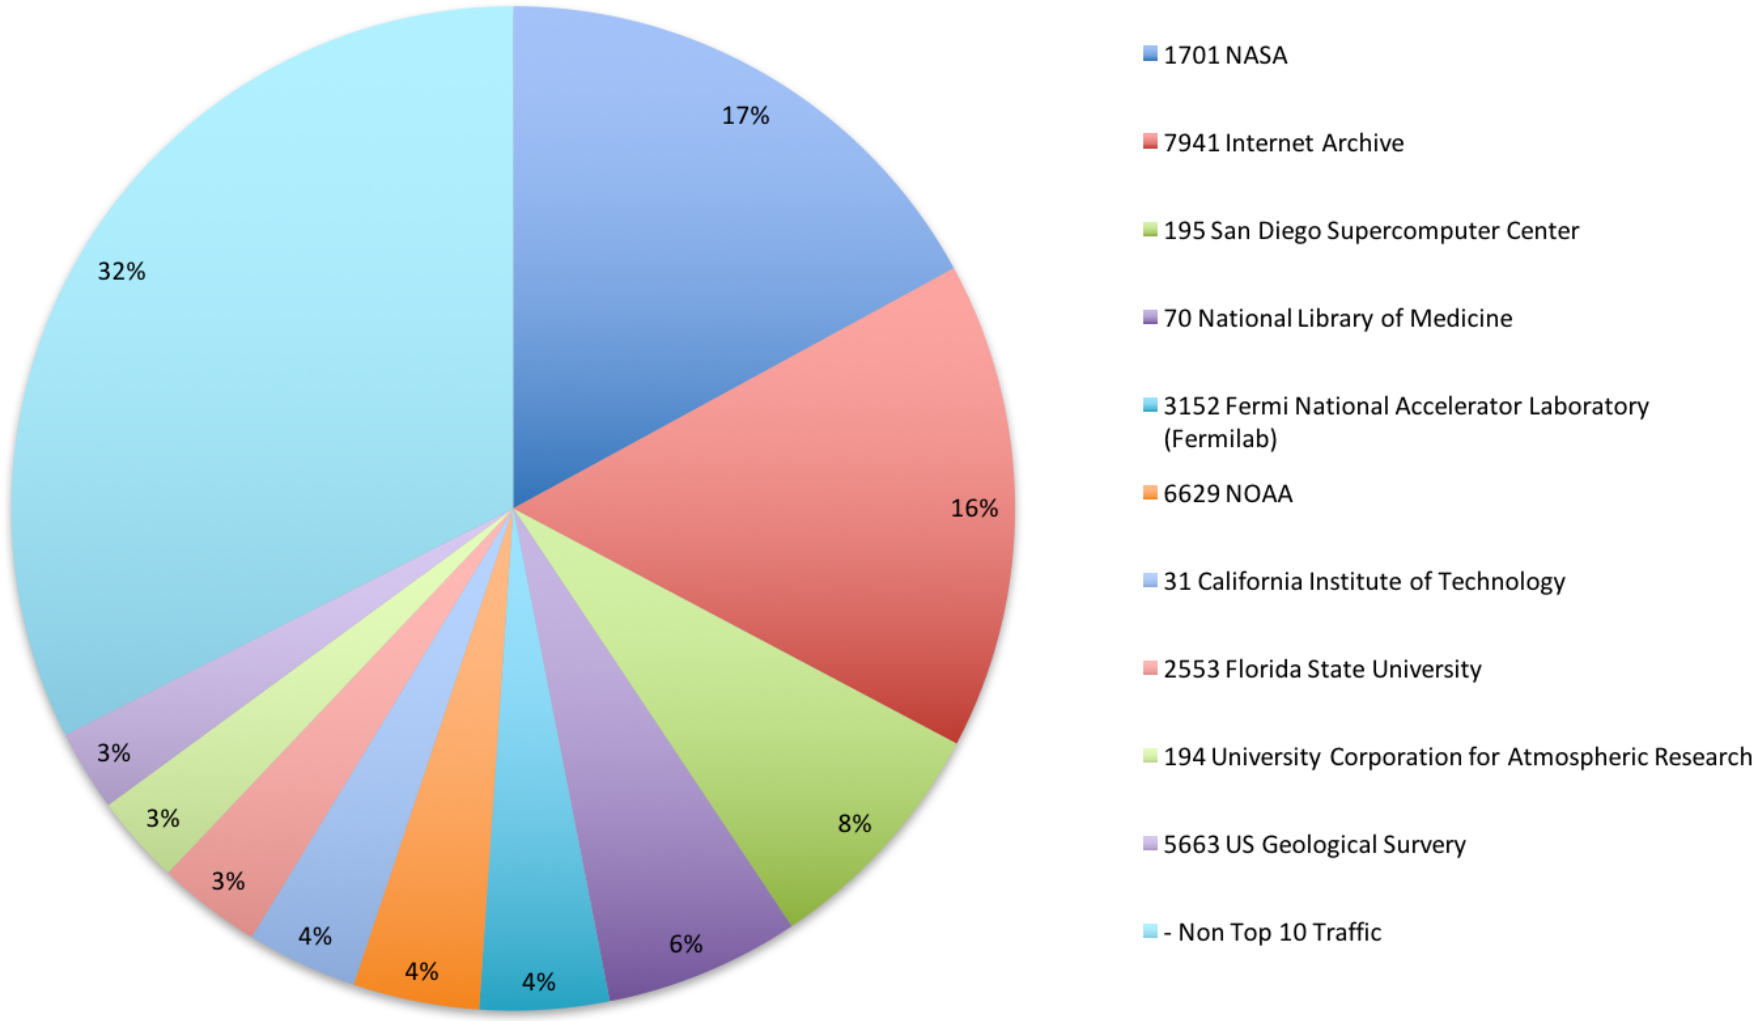

TransPAC4 Year 2 Quarter 1 Top Ten Talkers by Destination AS Outbound

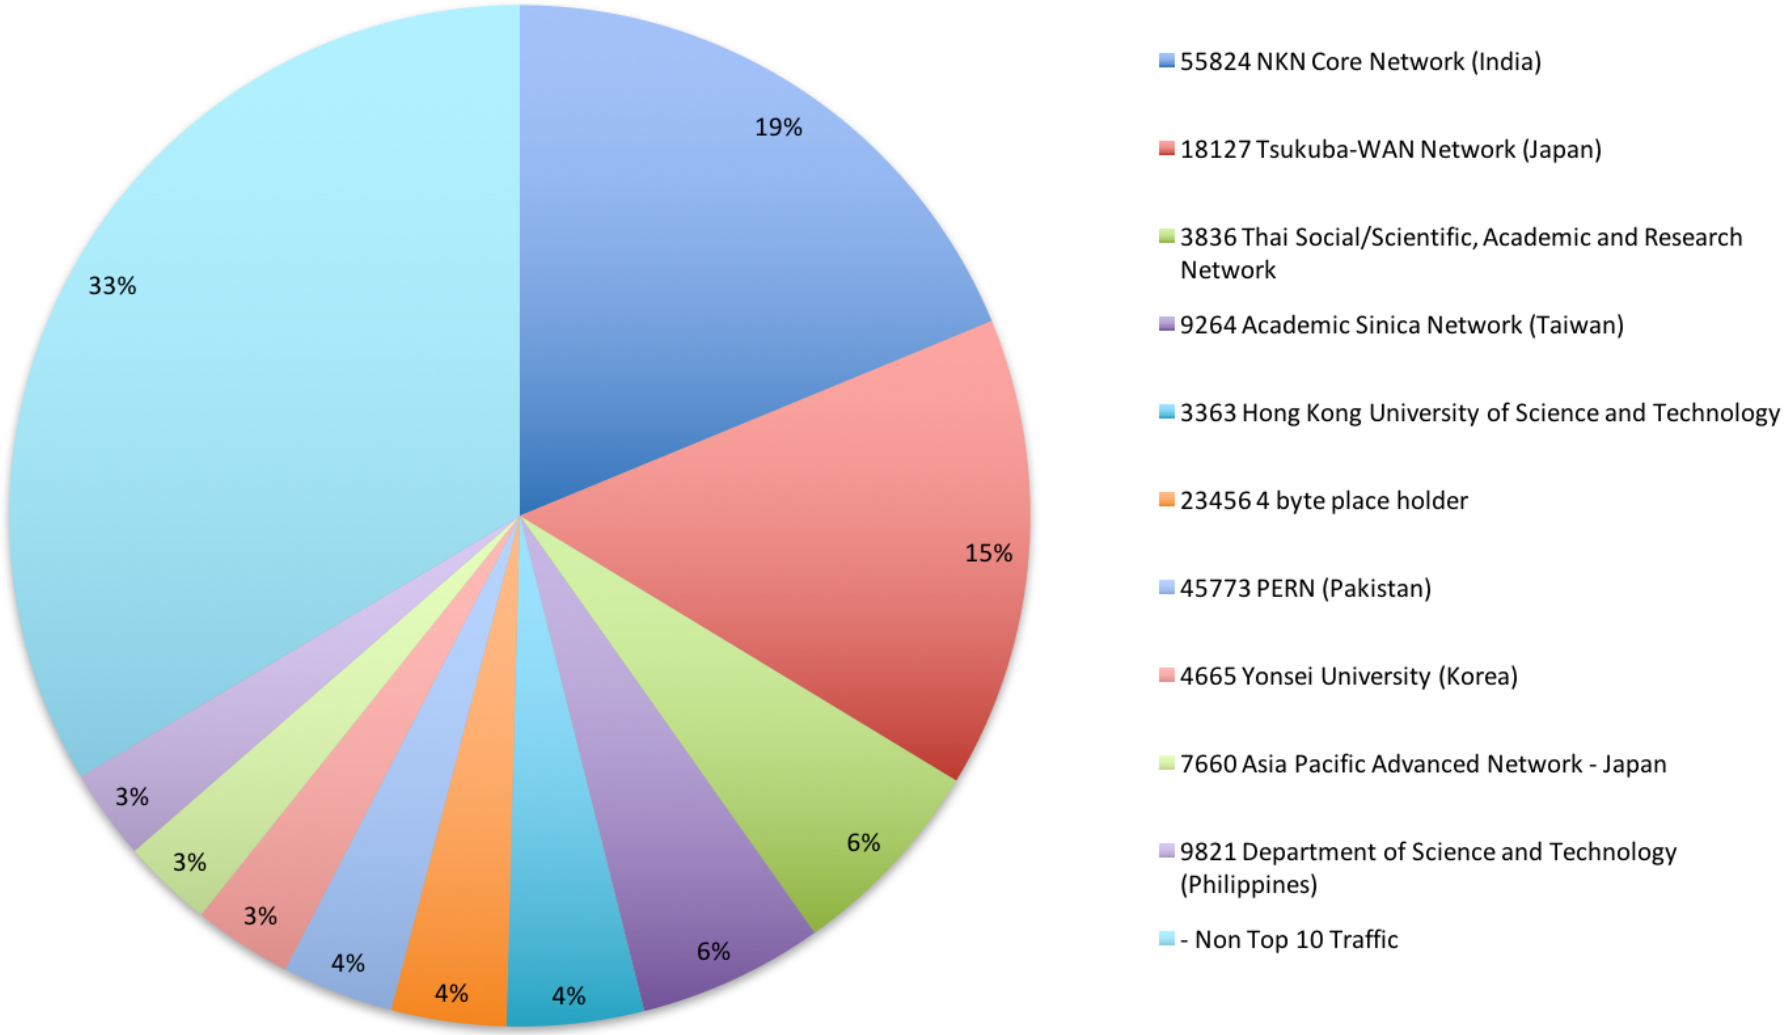

TransPAC4 Year 2 Quarter 1 Top Ten Talkers by Source AS Inbound

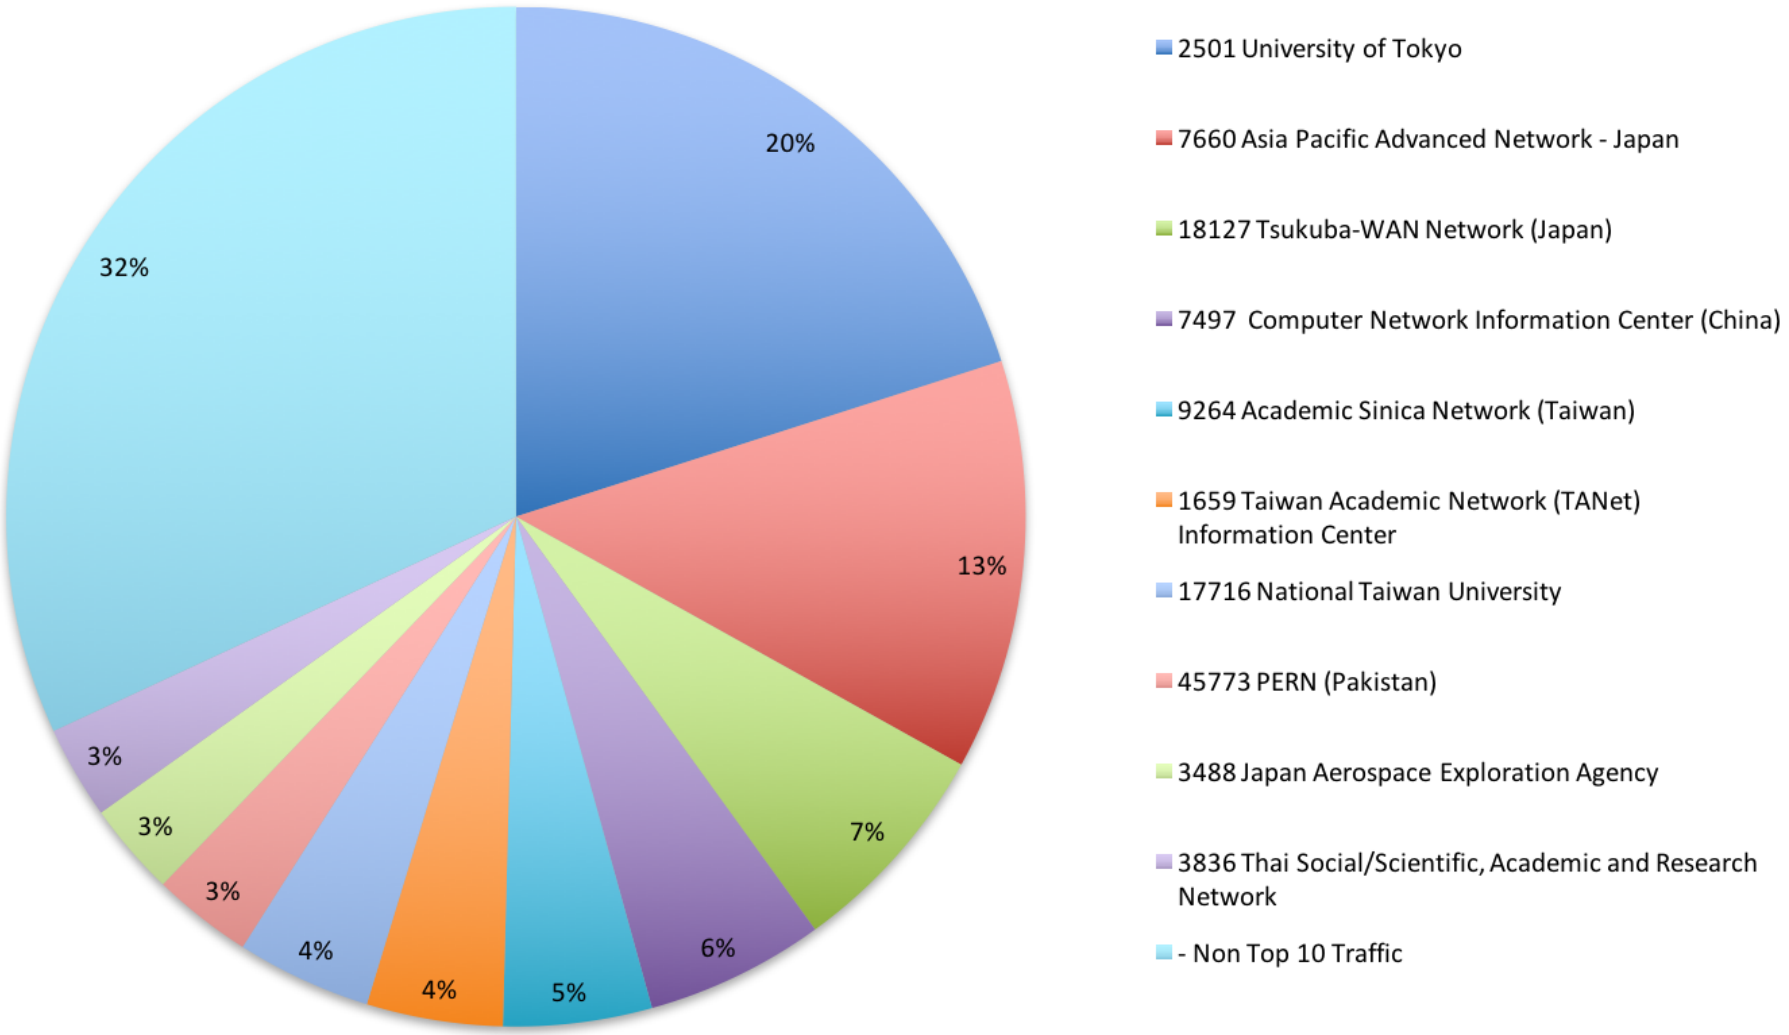

TransPAC4 Year 2 Quarter 1 Top Ten Talkers by Destination AS Inbound

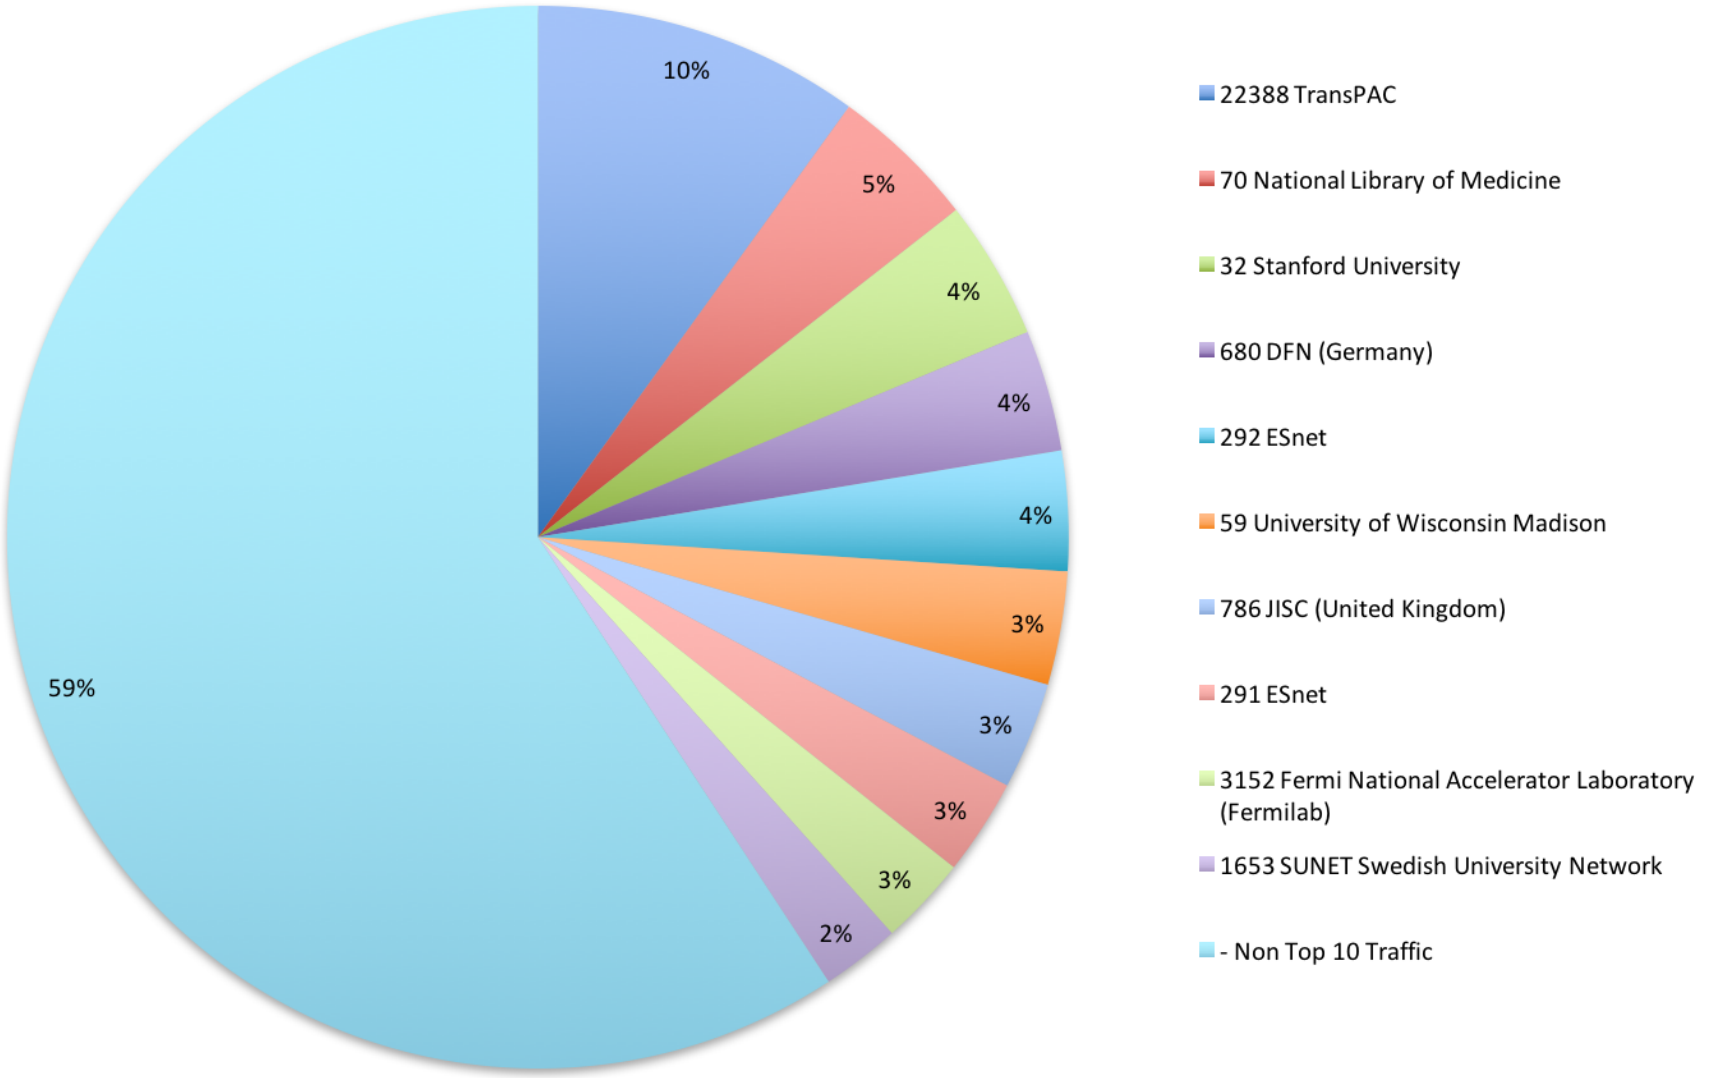

High Level Analysis

- Outbound
  - ~1/3 is geoscience data
- Inbound
  - 10-15% is EU pass through traffic
- Asian end points include
  - Japan, China, Taiwan, Pakistan, Hong Kong, India, Singapore, Thailand, Korea, etc

e-Science


e-Science Activity in Asia Pacific I
View details | Export
11:00 - 12:30
Room: BHSS, Conf. Room 2
Location: Academia Sinica
Convener: Dr. Gergely SIPOS
Contributions
11:00 eScience Activities in Japan
11:10 eScience Activities in China
11:20 eScience Activities in Korea
11:30 eScience Activities in Taiwan
11:40 e-Science Activities in Mongolian Academy of Sciences
11:50 Panel Discussion
View contribution list
e-Science Activities in Asia Pacific II
View details | Export
11:00 - 12:30
BHSS, Conf. Room 2
Academia Sinica
Dr. Ludek MATYSKA
Contributions
11:00 eScience Activities in Thailand
11:10 eScience Activities in Indonesia
11:20 eScience Activities in Malaysia
11:30 eScience Activities in Vietnam
11:40 eScience Activities in Philippine
11:50 Panel Discussion
View contribution list
e-Science Activities in Asia Pacific III
View details | Export
11:00 - 12:30
Room: BHSS, Conf. Room 2
Academia Sinica
Dr. Alberto MASONI
EU-India cooperation on e-Infrastructures: from EU-IndiaGrid to e-INIT project
eScience Activities in India
eScience Activities in Australia
eScience Activities in Singapore (remote presentation)
eScience Activity in Nepal
Panel Discussion
contribution list

Japan

- We peer with APAN-JP (AS 7660)
- SINET (AS 2907) has their own trans Pacific connectivity
  - We still see some data – 1.8 TB
  - JAXA
  - Lots of networks behind
  - Ministry of Agriculture, Forestry and Fisheries Research Network (AS18125) – 13 TB
    - Top transfer from IU Biology dept – 1 TB
- NICT/JGN-X(AS 9355)
  - Transferred 3.7 TB
  - Still more geoscience

China

- CSTNet - AS 7497 – 36TB
  - 4 TB worth of Biology data
  - Other intuitions behind
- CERNET – AS 4538
  - CERNET has their own 10G
  - 34 TB carried
  - Most of it Genomics

Korea

- KISTI backbone – AS 17597
  - Lots if institutions behind
- KISTI (AS 1237)
  - 4TB data transferred
  - Some of which goes to somewhere behind NORDUnet

Taiwan

- **Has its own connectivity**
- **Academia Sinica, TANET and University of Tawain still show up as top talkers**

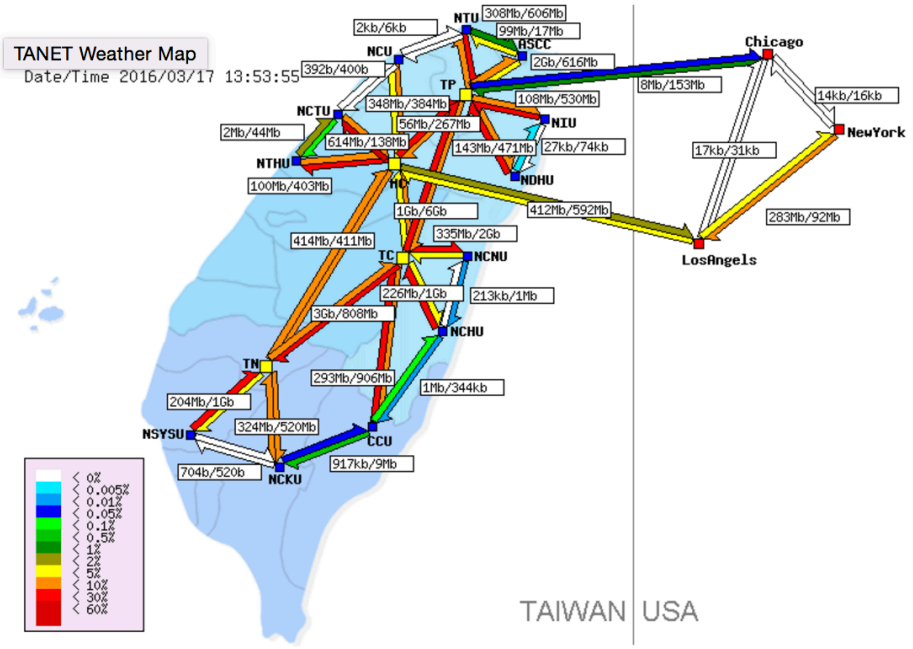

Taiwan

- Academia Sinica (AS 9264) – 95TB
  - Climate data
- TANET (AS 1659) – 22.5TB
  - Also lots of climate data
  - Traffic from Europe
- National Taiwan University (AS 17716) – 10.8 TB
  - >2TB HEP data to CERN, Purdue

Mongolia

- Mongolian Academy of Sciences
  - Appears to use the commodity internet

Thailand

- ThaiREN (AS 24475)
- Thai Social/Scientific, Academic and Research Network (AS 3836)
  - National Science and Technology Development Agency (AS 38296)
  - UNINET-TH (AS 4621) – 30.6 TB
    - NASA, Goddard Space Flight Center

Indonesia

- INHERENT (AS 18007)
- Lots of universities
- Institute of Technology Bandung (AS 4796)
  - 2 TB, much of it climate data

Malaysia

- MYREN (AS 24514)
- 12TB transferred
- Other networks behind MYREN
- UPM (AS 38868) – 12TB
- UTM (AS 133014) – 4 byte AS
  - Academic Grid Malaysia Science Gateway
- USM (AS 24090) – 1.8TB

Vietnam

- VAST
- NASTI
- VinaREN
- No traffic detected, but:

- Aug 2014: Cyberinfrastructure and Water Resources in the Lower Mekong Region
- http://internationalnetworking.iu.edu/archives/LMI/

Philippines

- Philippines Department of Science and Technology (AS 9821)
  - 33 TB
  - University of the Philippines System Network (AS 132483)
  - University of the Philippines Diliman (AS 23862) – 3TB

India

- Networking in India is very complex
- NKN (AS 55824/55847)
  - 277TB data transferred!
    - LIGO
  - 22 different ASNs behind NKN, 13 4 byte
  - CDAC (AS 55433) - 1.5TB

Australia

- AARNet (AS7575) obviously has its own network
  - SXTransport
  - Still see a lot of data going across TransPAC
  - Transiting data from Australia to Asia ~1TB
    - Especially NCI->Academica Sinica
  - Environmental, IHEP-CN
- Other networks behind AARNet
- Might not be the most optimal situation

IRNC funded Backbones and Exchange Points

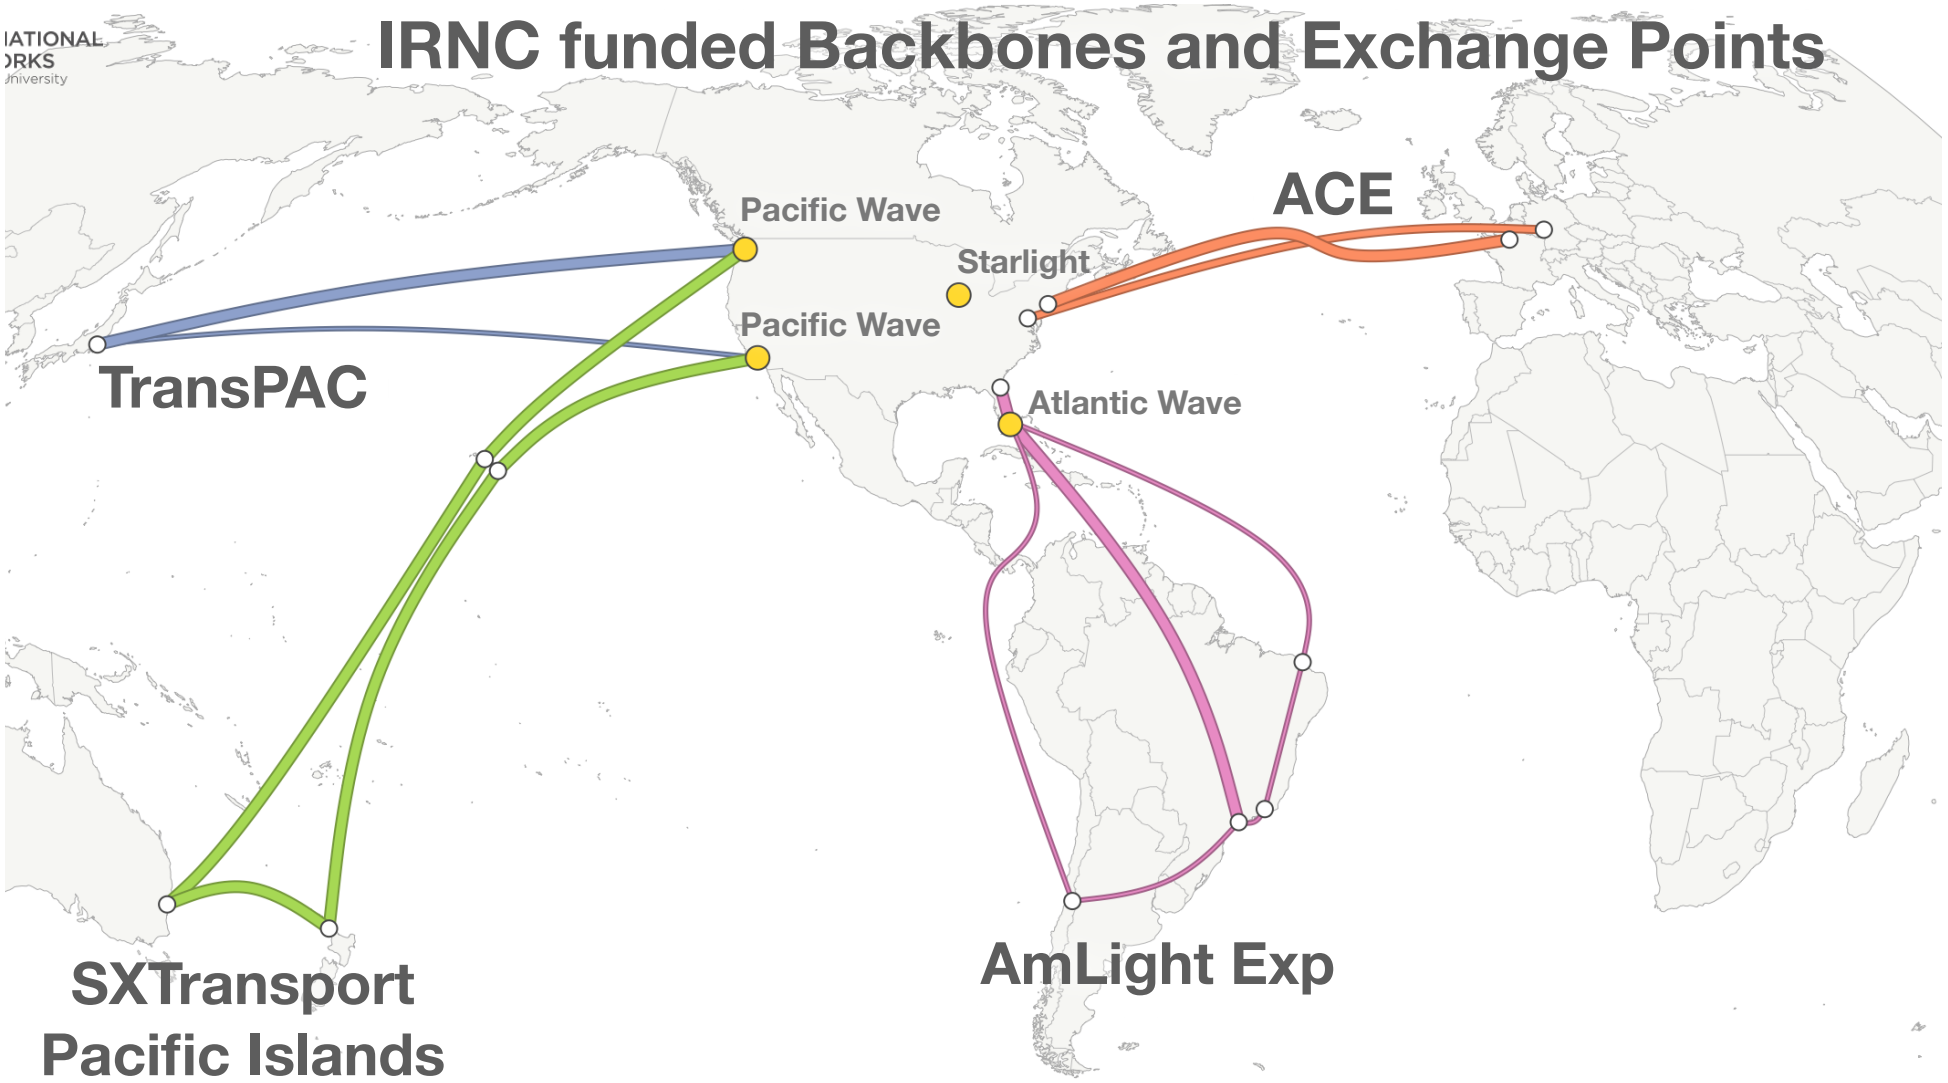

Singapore

- SINGAREN has its own 100G circuit to Los Angeles in partnership with Internet2
- Given that you would think we woudn't see much traffic, but we do, not from SingaREN (AS 23855) itself but networks behind it
- A*STAR (AS 23767) - 13TB transferred
- Other networks behind SINGAREN

Nepal

- **Connectivity from the US seems to be across the commercial Internet**

Science Engagement Specialist Predrag Radulovic

- Proactive interactions with end-user science community
- Identify large collaborations sharing data
  - Flow Data
  - End users at meetings
- Assist in running workshops to bring end users and network engineers together

End User Engagement

- Directly engage with user communities
- Focusing on end-to-end performance
- Application Advocacy Specialists will work closely with end users with network needs
  - including US branch campuses in Asia
- Why do we need this?

CrossConnects
WORKSHOP SERIES

Where Science & Technology Meet

Save the Date

CrossConnects Workshop: Improving Data Mobility & Management for International Bioinformatics

When: April 12-13, 2016
Where: Lawrence Berkeley National Lab, Berkeley, CA

The CrossConnects Workshop aims to bring together leaders in the bioinformatics, computing, and networking communities to discuss the resources, partners, and tools needed to support high performance data transfers, distributed data analysis and global collaboration in precision medicine, precision agriculture and their relevant ties to human and plant microbiomic and metagenomic research.

International participation is strongly encouraged for this workshop that aims to engage researchers and practitioners in precision medicine and metagenomics, data managers, networking experts, and software engineers.

More information at: http://www.es.net/crossconnects-workshop/
Questions? Email: cc-bio-pc@lists.lbl.gov

Hosted by:

INDIA

INTERNATIONAL NETWORKS
At Indiana University

TransPAC plans additional circuits

- **Exploring alternative options away from our traditional paths**
  - **Guam to ?**
- **Peer in Tokyo and/or Seattle with all other willing R&E nets or peering points**


Seattle
Beijing
Seoul
LA
Hong Kong
Taipei
Okinawa
Manila
Guam
Vietnam
Kuala Lumpur
Singapore
Indonesia
Auckland

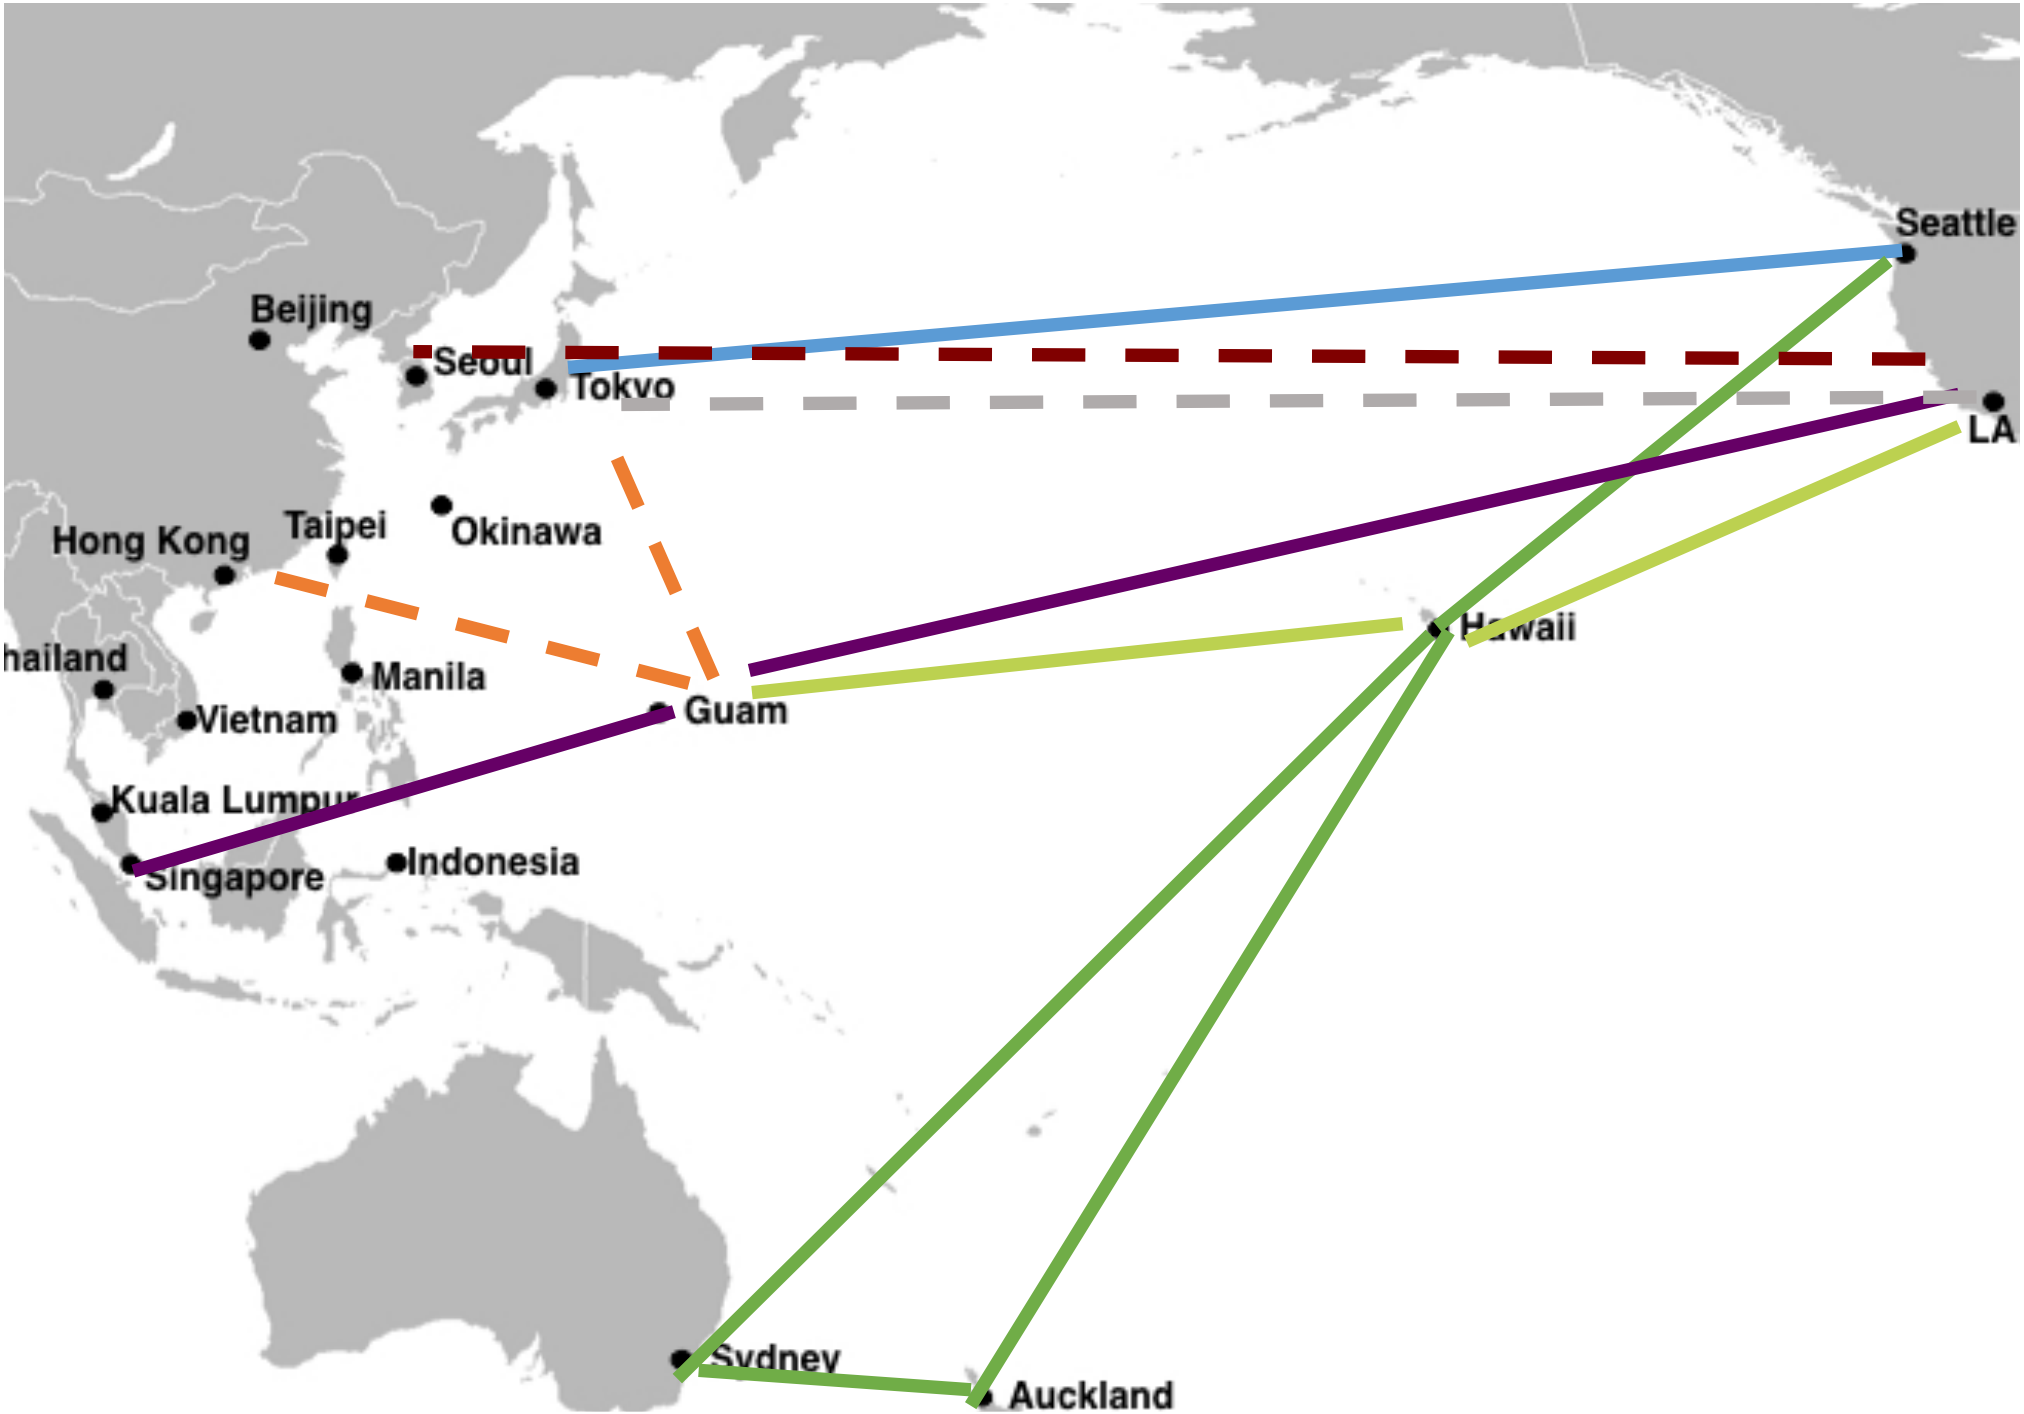
Seattle
Beijing
Seoul
Tokyo
LA
Taipei
Okinawa
Hong Kong
Manila
Guam
Vietnam
Kuala Lumpur
Singapore
Indonesia
Sydney
Auckland

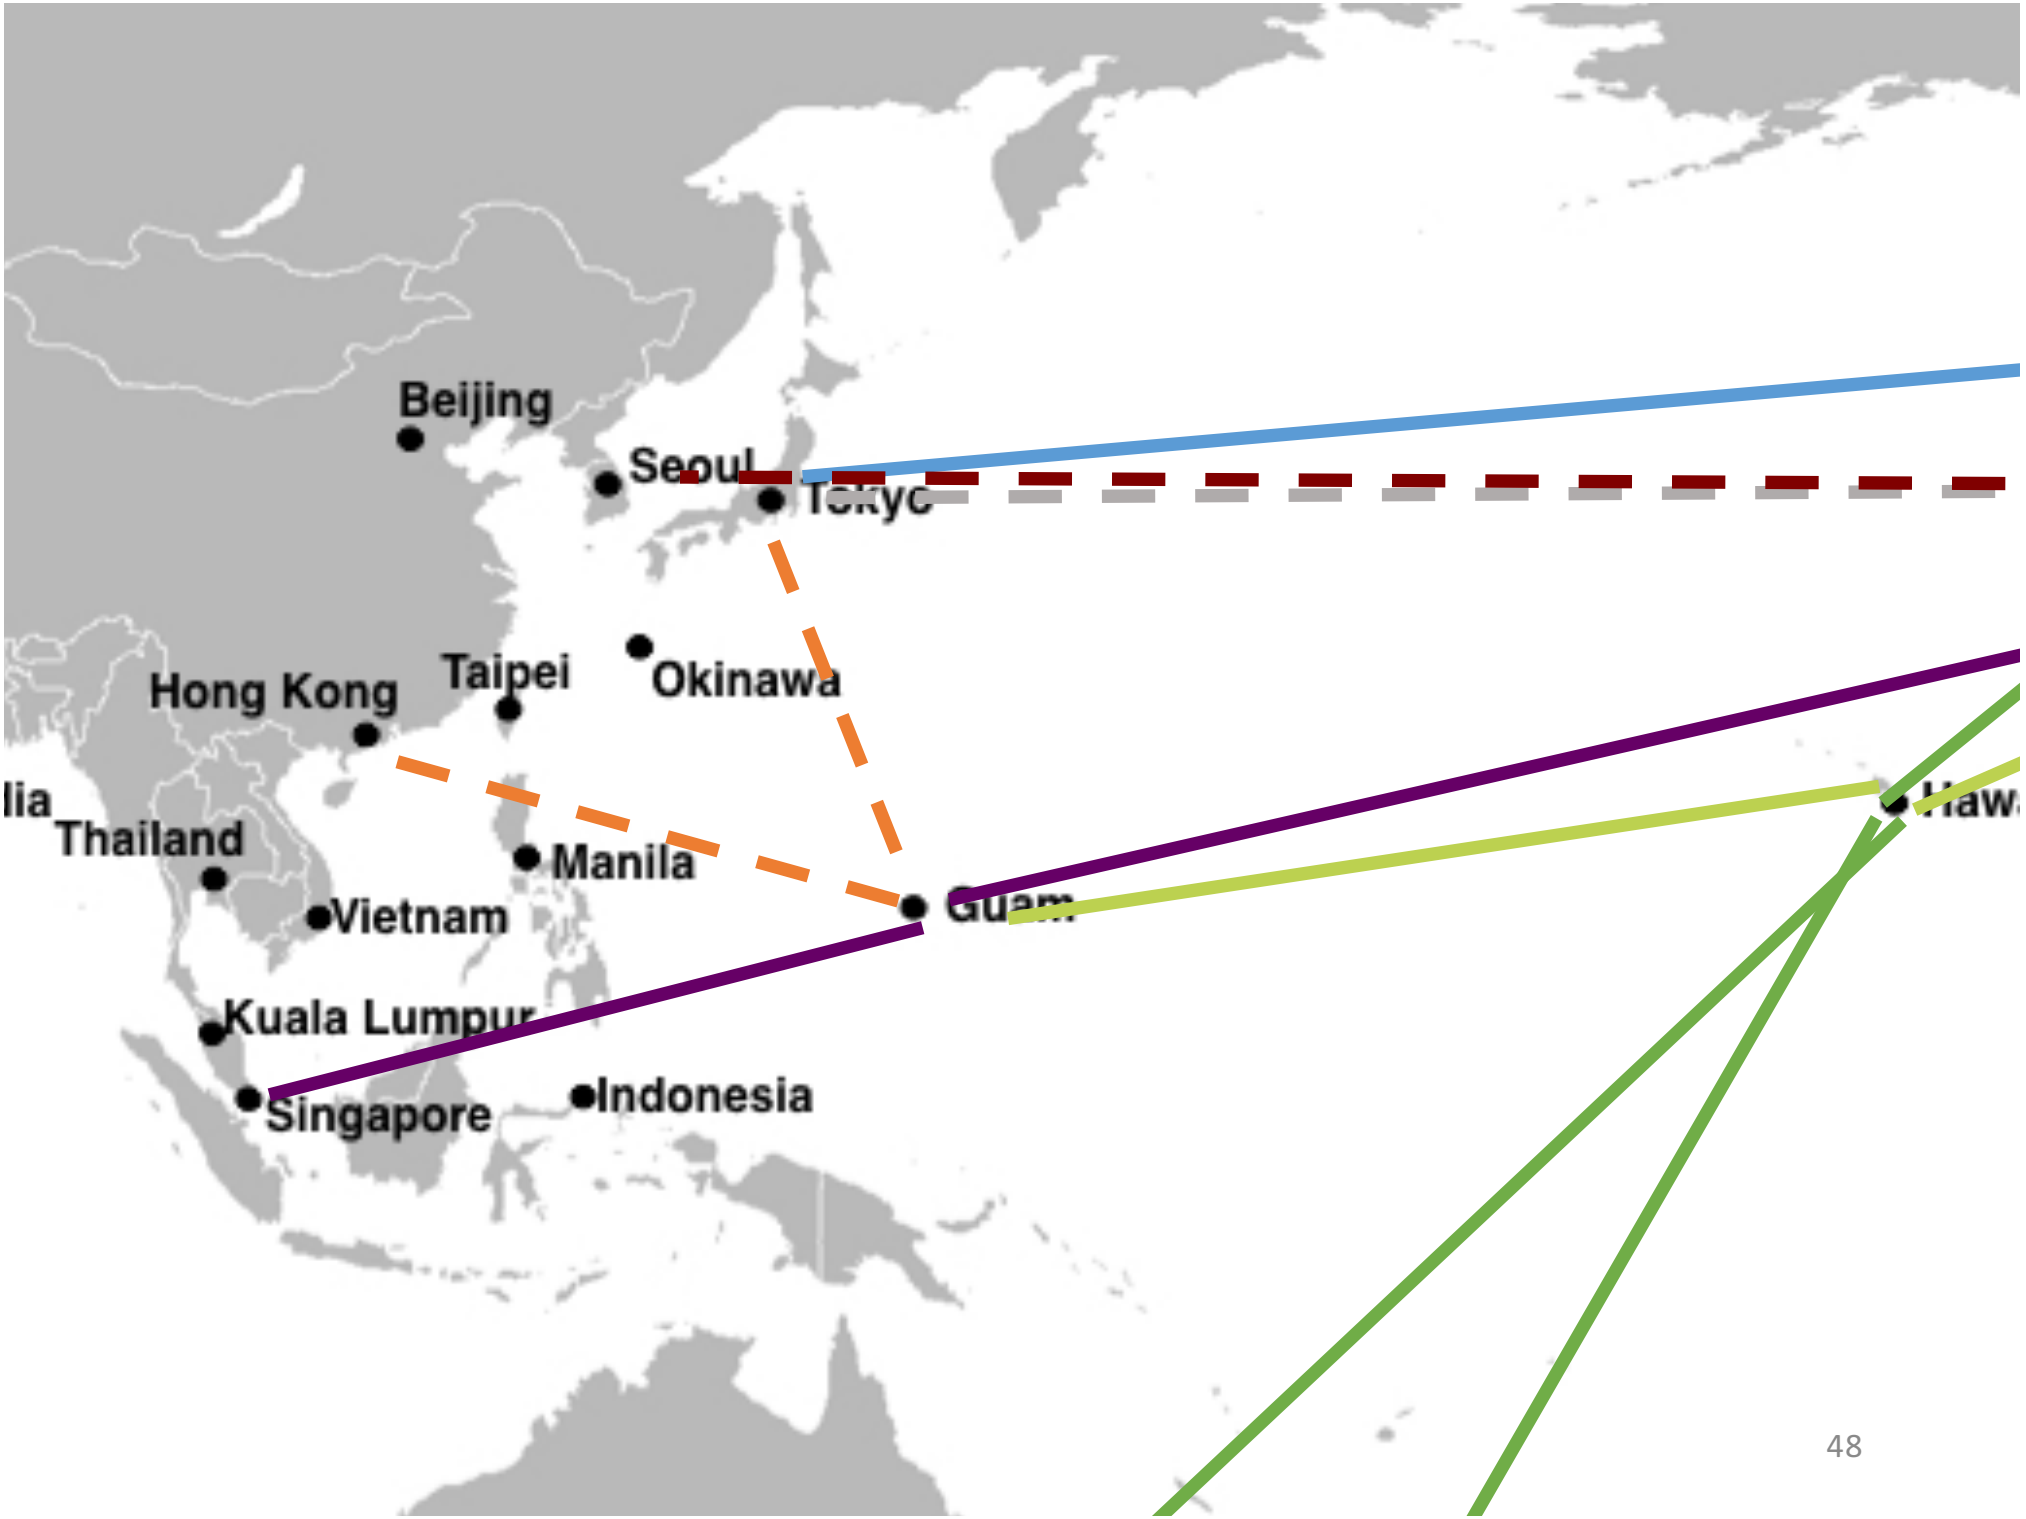
Beijing
Seoul
Tokyo
Hong Kong
Taipei
Okinawa
Thailand
Manila
Vietnam
Guam
Kuala Lumpur
Singapore
Indonesia

TEIN

Connecting Asia and Europe's Research and Education Communities

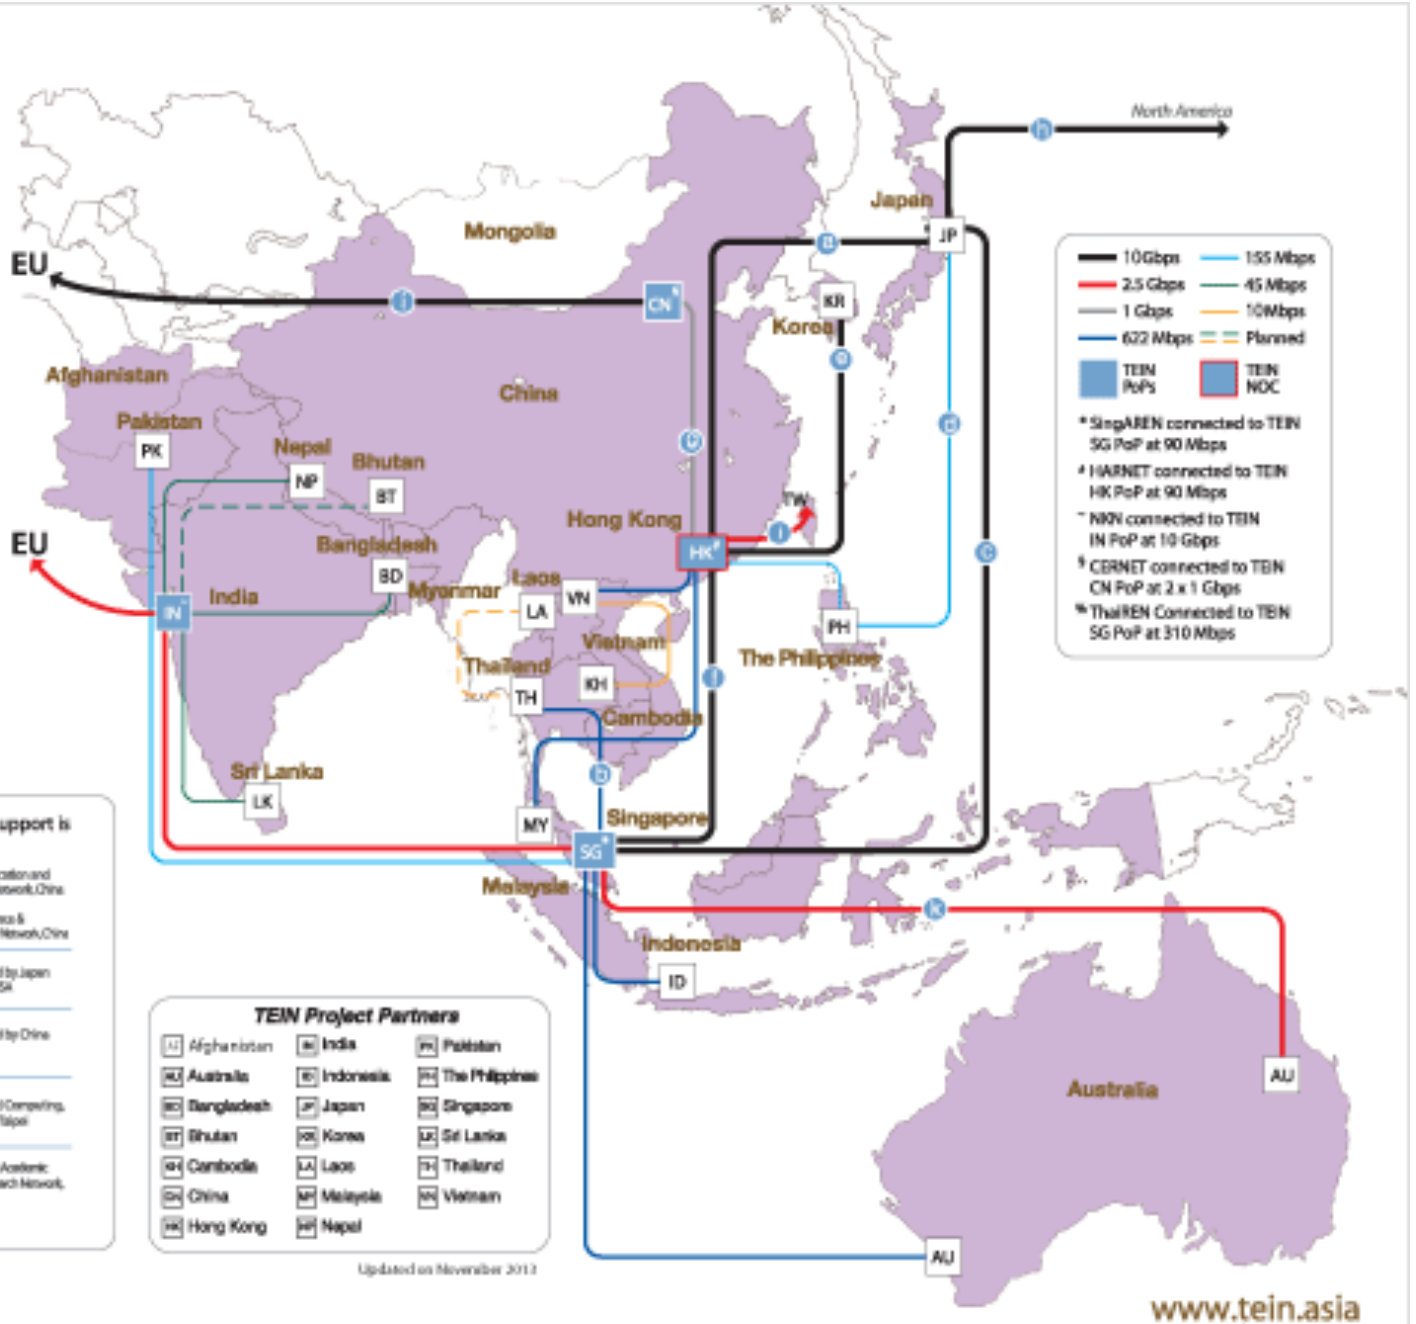

The following links are fully financed by the link owners whose support is gratefully acknowledged.

- a, b, c, d, e, f: NICT, National Institute of Information and Communications, Japan
- b: Thailand Research and Education Network, Thailand
- c: NII, National Institute of Informatics, Japan
- d: MAFFIN, Ministry of Agriculture, Forestry and Fisheries Research Information Network, Japan
- e, f: NIA, National Information Society Agency, South Korea
- g: China Education and Research Network, China; China Science & Technology Network, China
- h: TransPAC, co-funded by Japan and the USA
- i: ORIENT plus, co-funded by China and EU
- j: ASGC, Academia Sinica Grid Computing, Republic of Chinese Taipei
- k: AARNet, Australia Academic and Research Network, Australia

TEIN Project Partners

AF Afghanistan	IN India	PK Pakistan
AU Australia	ID Indonesia	PH The Philippines
BD Bangladesh	JP Japan	SG Singapore
BT Bhutan	KR Korea	LK Sri Lanka
KH Cambodia	LA Laos	TH Thailand
CN China	MY Malaysia	VN Vietnam
HK Hong Kong	NP Nepal	

Updated on November 2013

www.tein.asia

This map has been produced with the financial assistance of the European Union. The contents of this document are the sole responsibility of TEIN*CC and can under no circumstances be regarded as reflecting the position of the European Union.

TEIN*CC COOPERATION CENTER
www.teincc.org

TEIN is co-funded by the European Commission through the Directorate-General for Development and Cooperation-EuropeAid

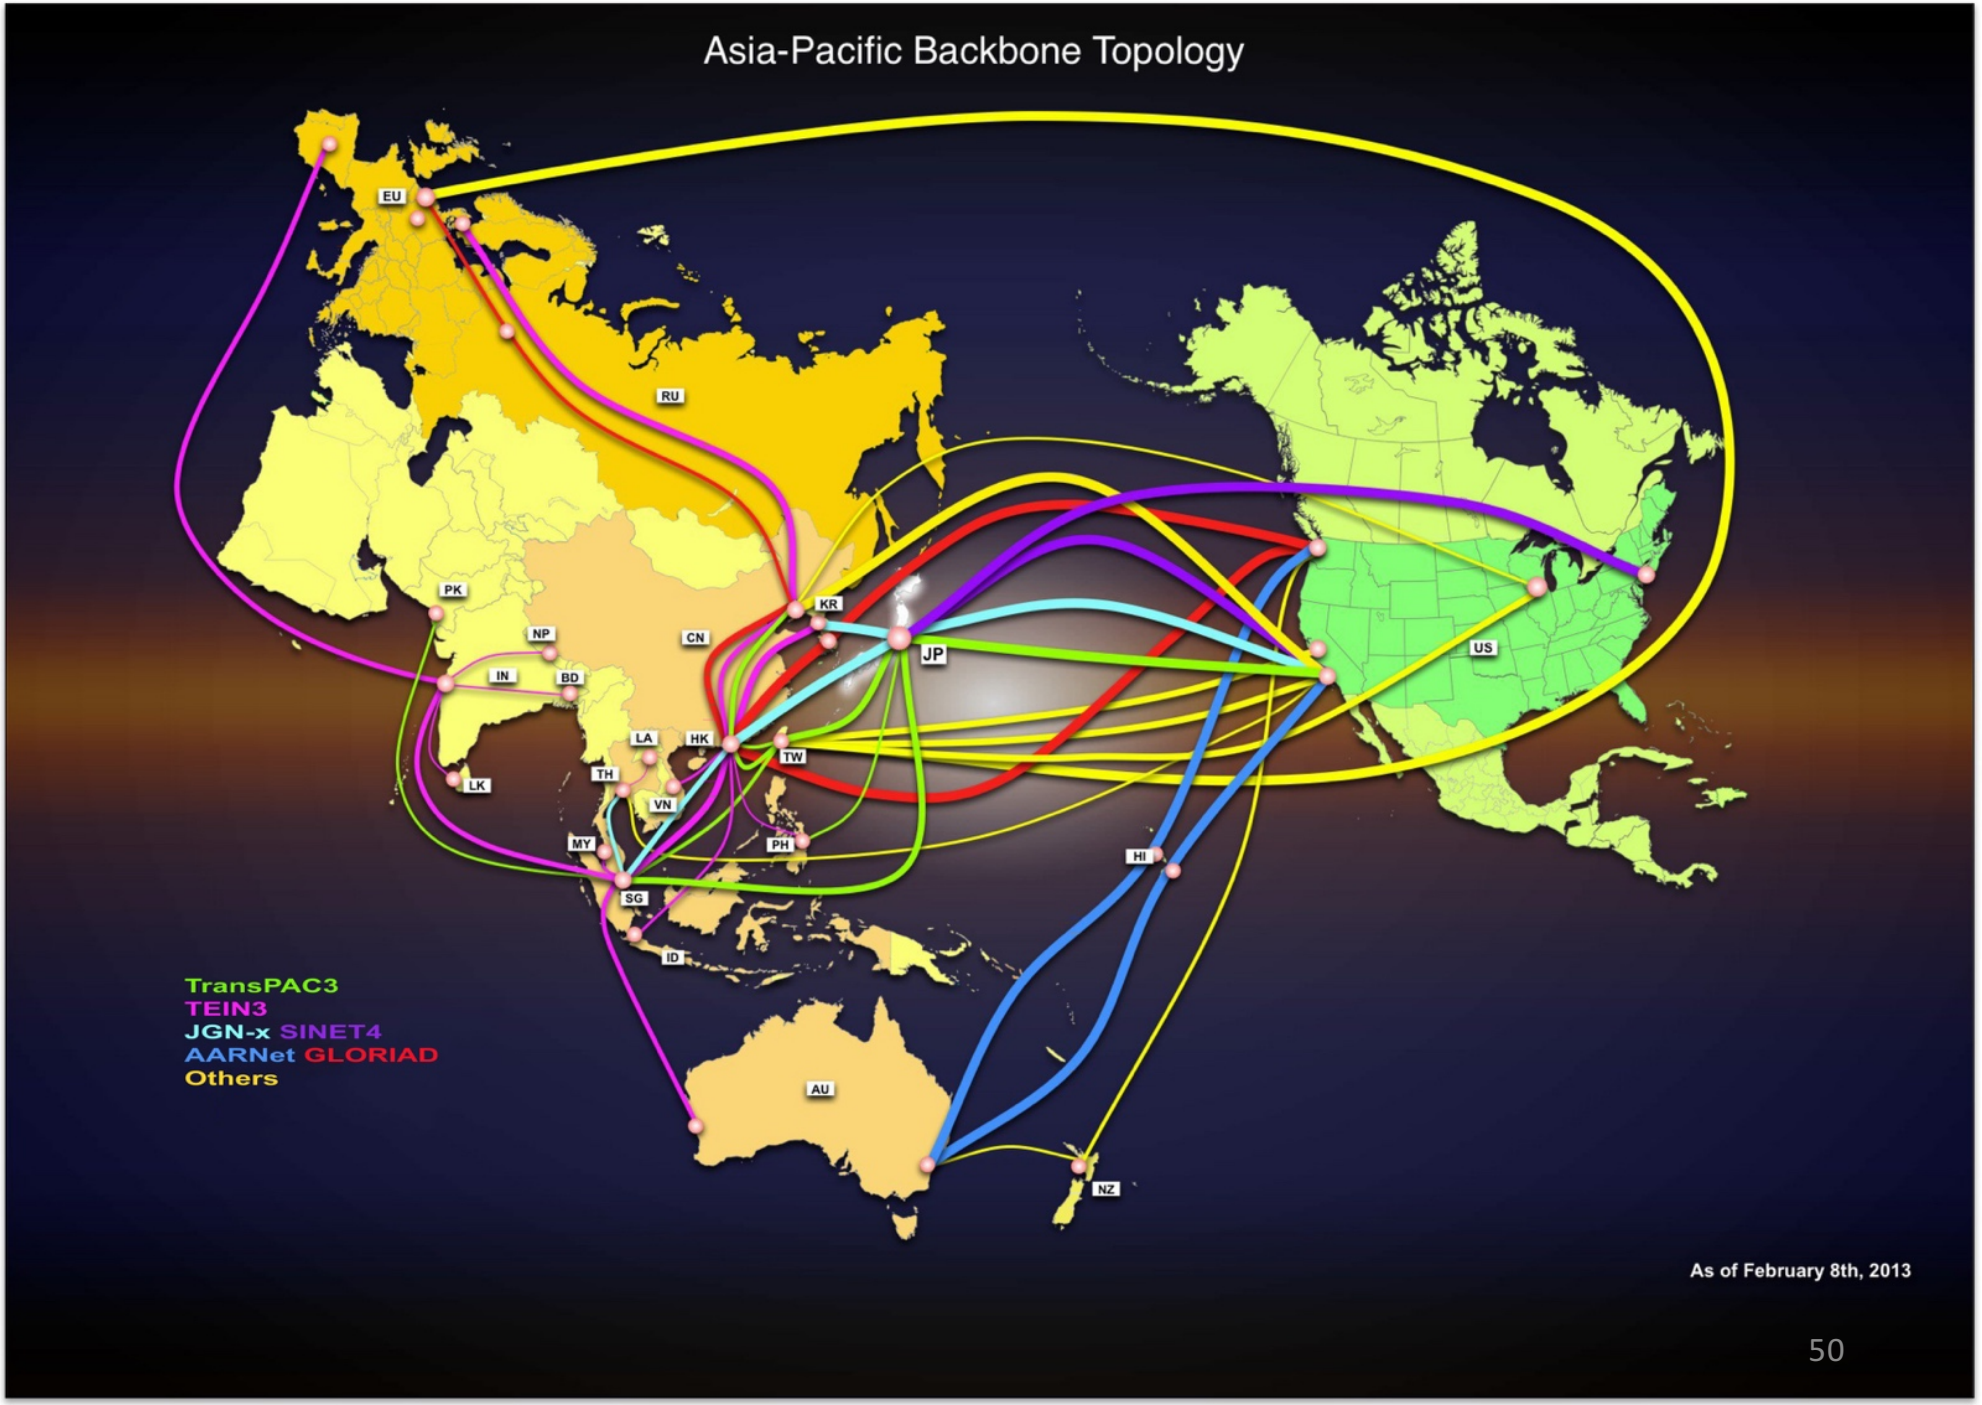
Asia-Pacific Backbone Topology
EU
RU
PK
NP
IN
BD
CN
KR
JP
US
LA
HK
TW
TH
LK
VN
MY
PH
HI
SG
ID
AU
NZ
TransPAC3
TEIN3
JGN-x SINET4
AARNet GLORIAD
Others
As of February 8th, 2013

perfSONAR in Emerging Regions

- Partnering with NSRC and ESnet
- Help expand the number of perfSONAR deployments was part of the work proposed
- Two-pronged approach
  - Training for network engineers
    - Hands on exercises
    - Development of additional materials including videos
  - Equipment for perfSONAR nodes, including installation assistance

IRNC Measurement: NetSage

- Joint project with
  - Indiana University
  - University of California at Davis
  - ESnet
  - University of Hawai'i Manoa
- Open, privacy-aware, network measurement, analysis, and visualization service
- Designed to address the needs of today's international networks
- Coordinated approach across backbones and exchange point

IRNC NOC also has Performance Engagement Team (PET)

- Reactive assistance for network errors
- PET will
  - Identify path
  - Investigate with network contacts
  - Test with available measurement points
  - Resolve and confirm resolution
- Researchers and network engineers can initiate an investigation
- Issues will be tracked in a ticketing system -> creates accountability, data, and centralized contact tracking

IRNC NOC

- Single, centralized NOC for IRNC funded infrastructure projects
  - A 24x7x365 NOC
  - Focus on proactive response,
  - Rich and transparent operational reporting,
  - End-to-end performance engagement.
- GlobalNOC Service Desk as centralized first point of contact
  - End-to-end performance assessment/triage
  - Pass Tier II and Tier III issues to the awarded infrastructure projects' engineering personnel
  - Coordinating problem resolution lifecycle from start to completion.

Acknowledgements

- IN@IU is funded by
  - US NSF award #0962968 for TransPAC3
  - US NSF award #1450904 for TransPAC4
- The TransPAC PacificWave 100gb/s network fabric is provided by Pacific Northwest GigaPop
- Pacific Wave is partially funded by a NSF IRNC award to CENIC & Pacific NorthWest Gigapop

Questions/Comments?

- Take aways:
  - TransPAC4 is more than just circuits
  - Provides connectivity to the whole region

- Group Website:

  http://internationalnetworking.indiana.edu/

- Jennifer Schopf – jmschopf@indiana.edu
- Andrew Lee – leea@indiana.edu

3 NetSage use cases

- 1) Current traffic patterns across IRNC links, and the ability to anticipate growth trends for capacity-planning purposes;
- 2) The main sources and sinks of large, elephant flows to know where to focus outreach and training opportunities; and
- 3) Where packet loss is occurring, whether the cause is congestion or other issues, and what impact it has on end-to-end performance.

Year 1: Active and Passive Measurements on Backbones (1)

- TSDS Common Archive
  - Open Sourced, Shared with NOC
- PerfSONAR
  - Hook into existing testing frameworks
  - Pull data into common archive
  - Enable common queries across IRNC sites
- SNMP
  - IRNC NOC will be collecting this data
  - Data will be in common archive
  - Access via SNAPP tool
  - Common queries across all IRNC sites

Example NetSage visualization walkthrough

1

User enters natural-language-like query or selects from previous queries or popular queries

Query can be refined with pull down menu to help user figure out what they can query for

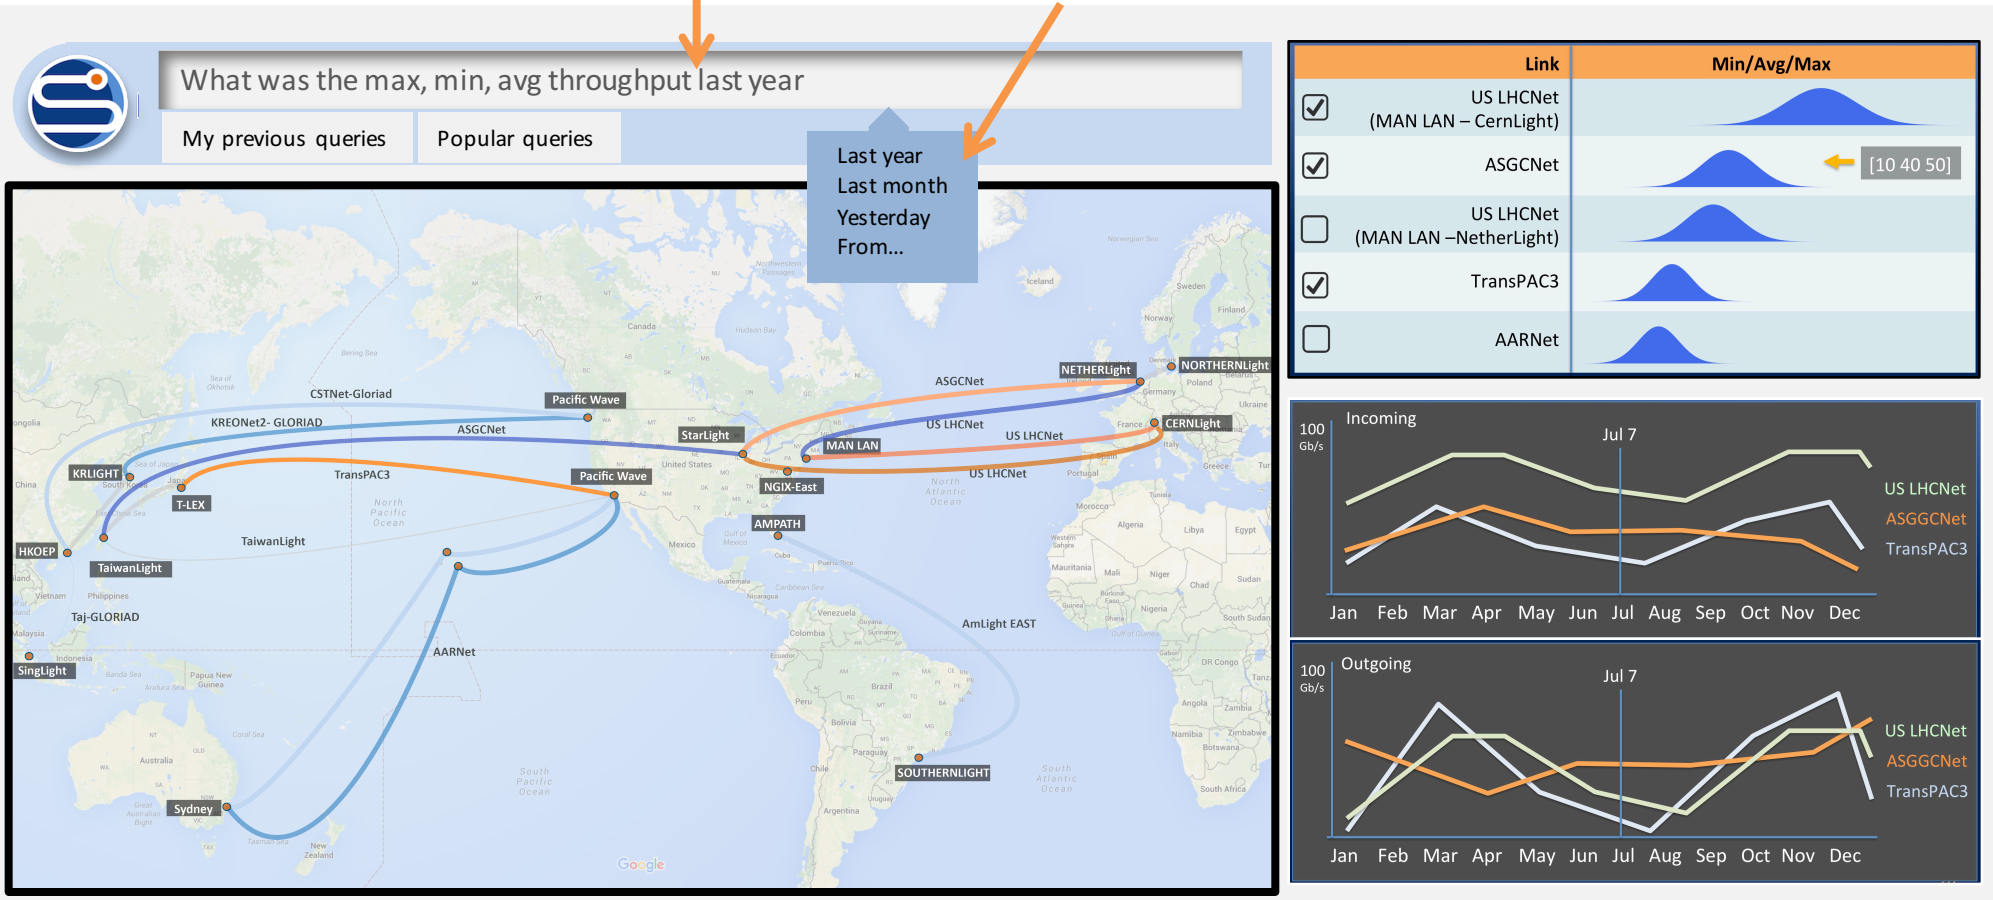

System always provides a high level contextual map of all IRNC links. Thickness of lines represent capacity. Color indicates current throughput & selected links for closer examination.

User can select multiple links or nodes to compare

Min, Avg, Max shown as stacked statistical curves

What was the max, min, avg throughput last year

My previous queries | Popular queries

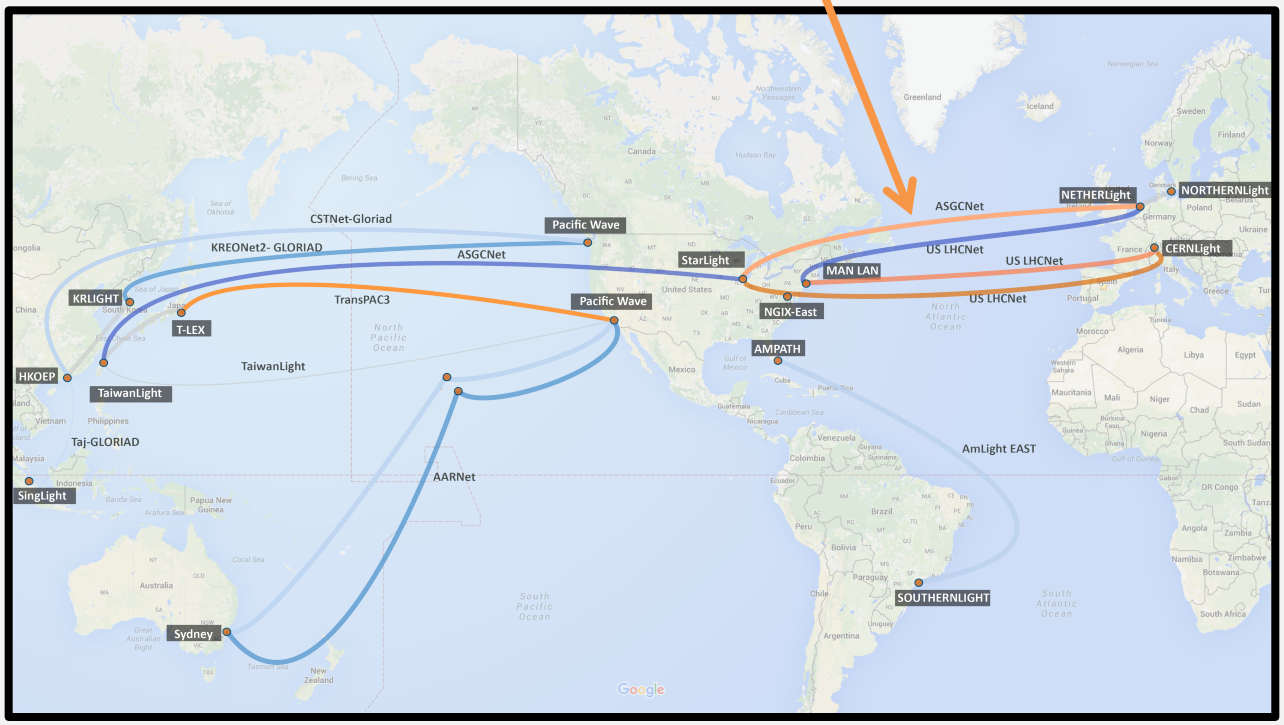

Link	Min/Avg/Max
☑ US LHCNet (MAN LAN – CernLight)	
☑ ASGCNet	[10 40 50]
☐ US LHCNet (MAN LAN –NetherLight)	
☑ TransPAC3	
☐ AARNet	

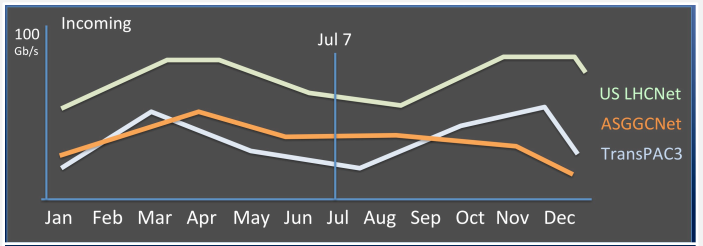

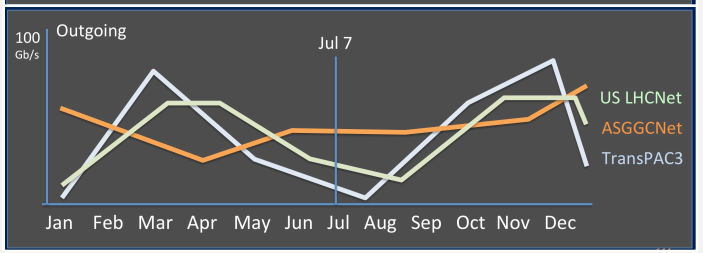

User can mark individual links to compare in greater detail

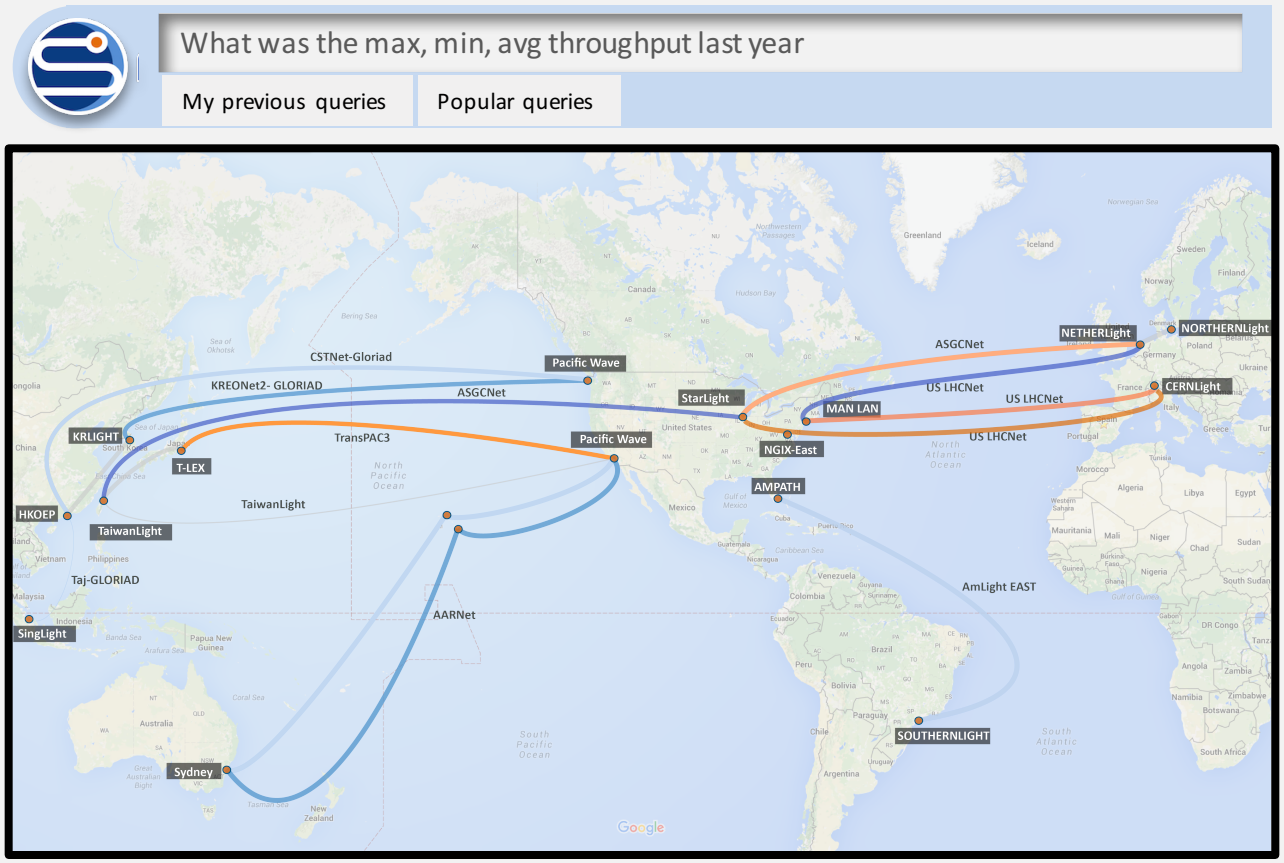

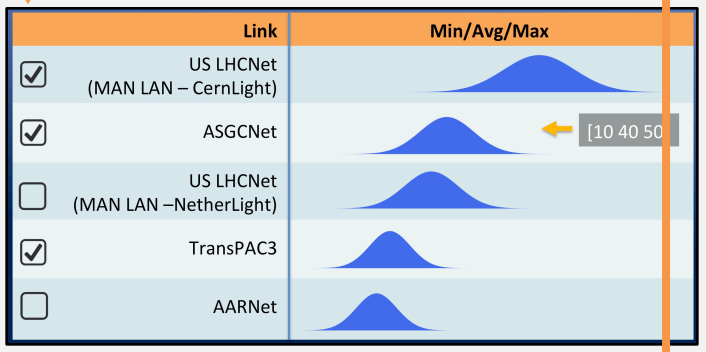

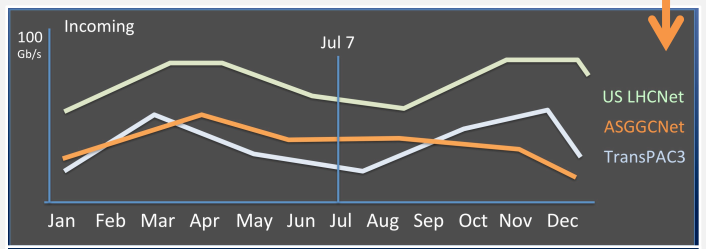

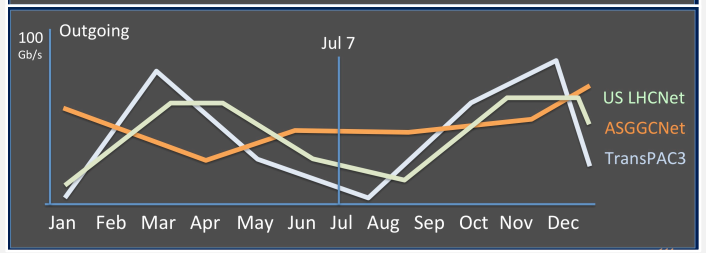

(showing incoming and outgoing throughput for the last year)

Any Questions so far?

IRNC funded Backbones and Exchange Points

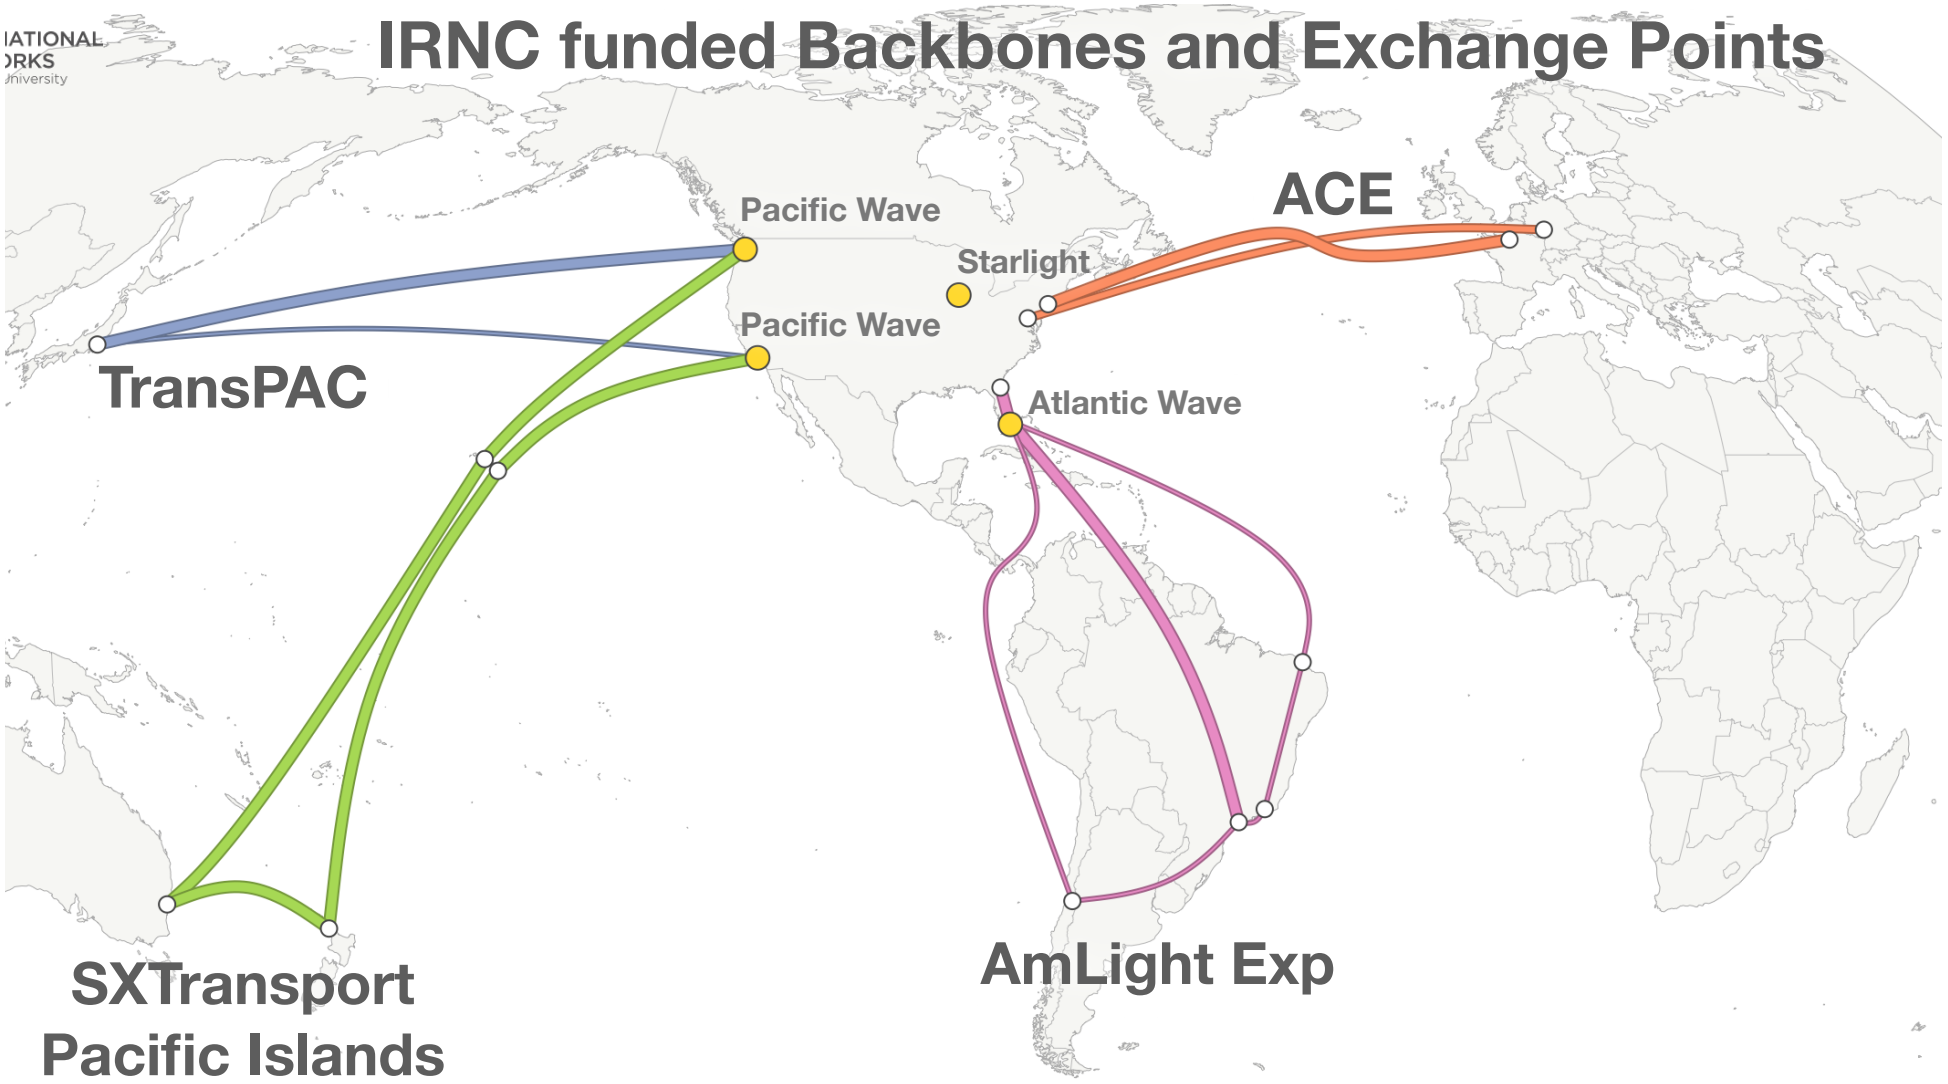

Transport

- Support for Layer 3 peering with TransPAC directly
- Multiple mechanisms for dynamic Layer 2 circuits
  - Joint with PacWave
- Peer at either Tata or WIDE/TREX

- Can we discuss peering and routing updates?

BGP Path Hinting

- **Inject additional metadata into the BGP routing table to help networks make more informed decisions about the best way to handle traffic**

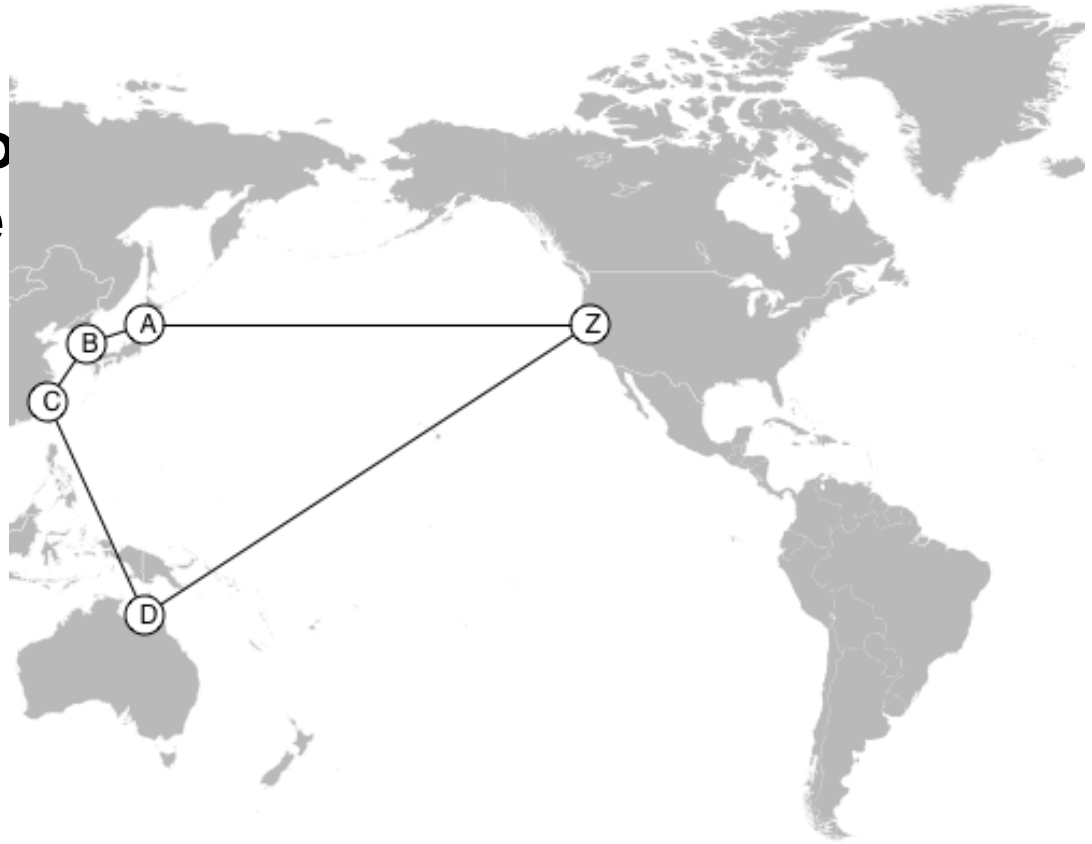

Next Steps

- Can we talk about peerings and routes?
- Can we talk about mutual backups and fail overs?