netsage

Open privacy-aware network measurement, analysis, and visualization service

Jennifer Schopf
International Networks, Indiana University

Sean Peisert, Brian Tierney
Esnet and UC Davis

Jason Leigh
Laboratory for Advanced Visualization & Applications,
University of Hawai'I Mānoa

Supported by the National Science Foundation

Acknowledgements

- NetSage is funded by US NSF award #1540933
- Joint project with
  - Indiana University
  - University of California at Davis
  - ESnet
  - University of Hawai'i Manoa

IRNC funded Backbones and Exchange Points

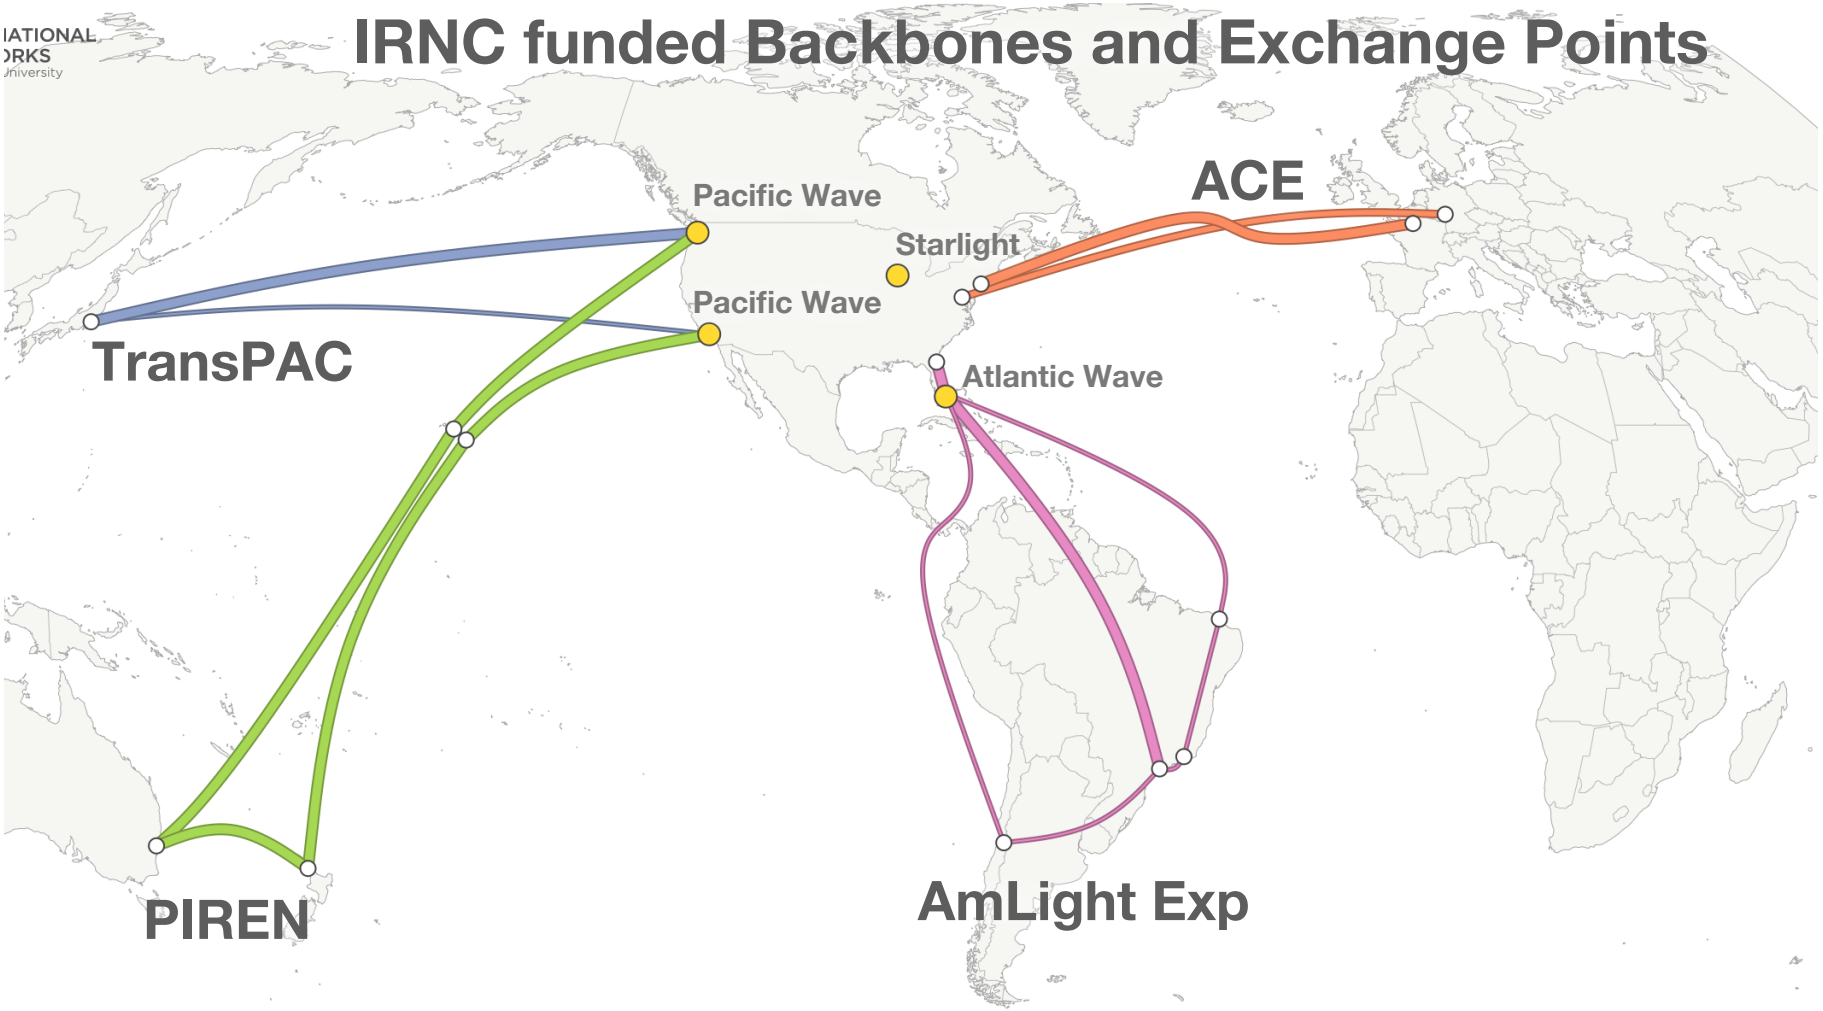

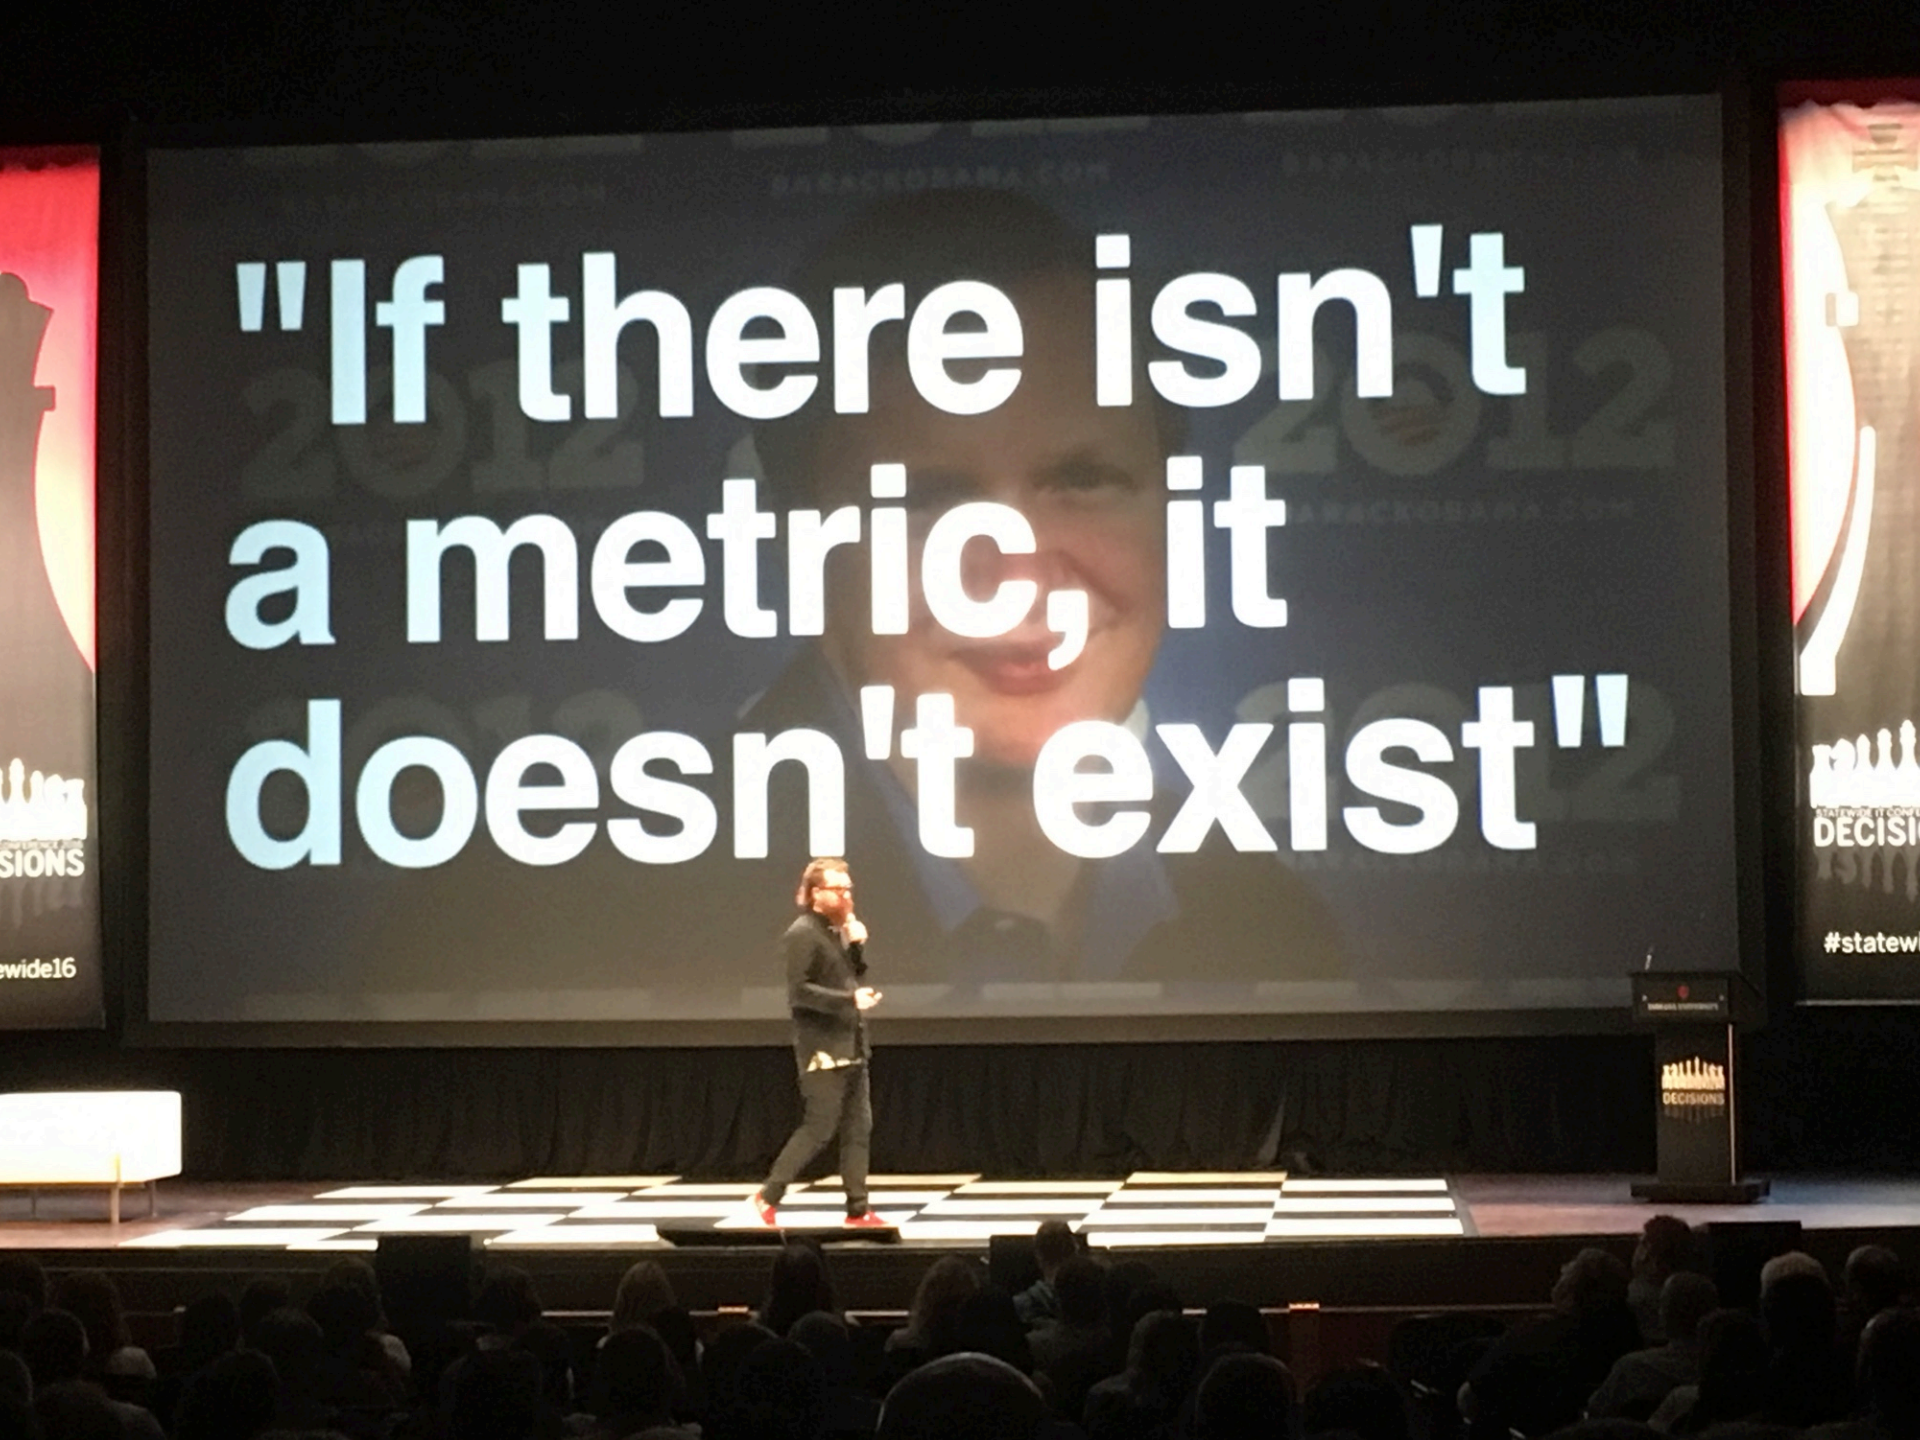
"If there isn't a metric, it doesn't exist"

IRNC Monitoring: NetSage

- Open, privacy-aware, network measurement, analysis, and visualization service
- Designed to address the needs of today's international networks
- Focus on work with production networks, not on novel research
- Tasked by NSF to coordinate monitoring work across the IRNC projects

IRNC Backbones and Exchange Points

- **Backbones:**
  - **TransPAC4 (Schopf, IU)**
  - **AmLIght ExP (Ibarra, FIU)**
  - **Pacific Islands Research and Education Networks (PIREN) (Lassner, U Hawaii)**
  - **America Connects to Europe (ACE) (Schopf, IU)**
  - **Network for European, American, and African Research (NEAAR) (Schopf, IU)**
- **Exchange Points**
  - **AtlanticWave (Ibarra, FIU)**
  - **StarLight (Mambretti, NU)**
  - **Pacific Wave(Fox, CENIC)**

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.

NetSage End users

- Project oversight (NSF, etc)
  - Congressional mandate that NSF projects show societal relevance
- IRNC NOC and other operators
- Project planning by backbone and exchange point operators
- Application engagement staff of IRNC projects

NetSage Measurement Architecture

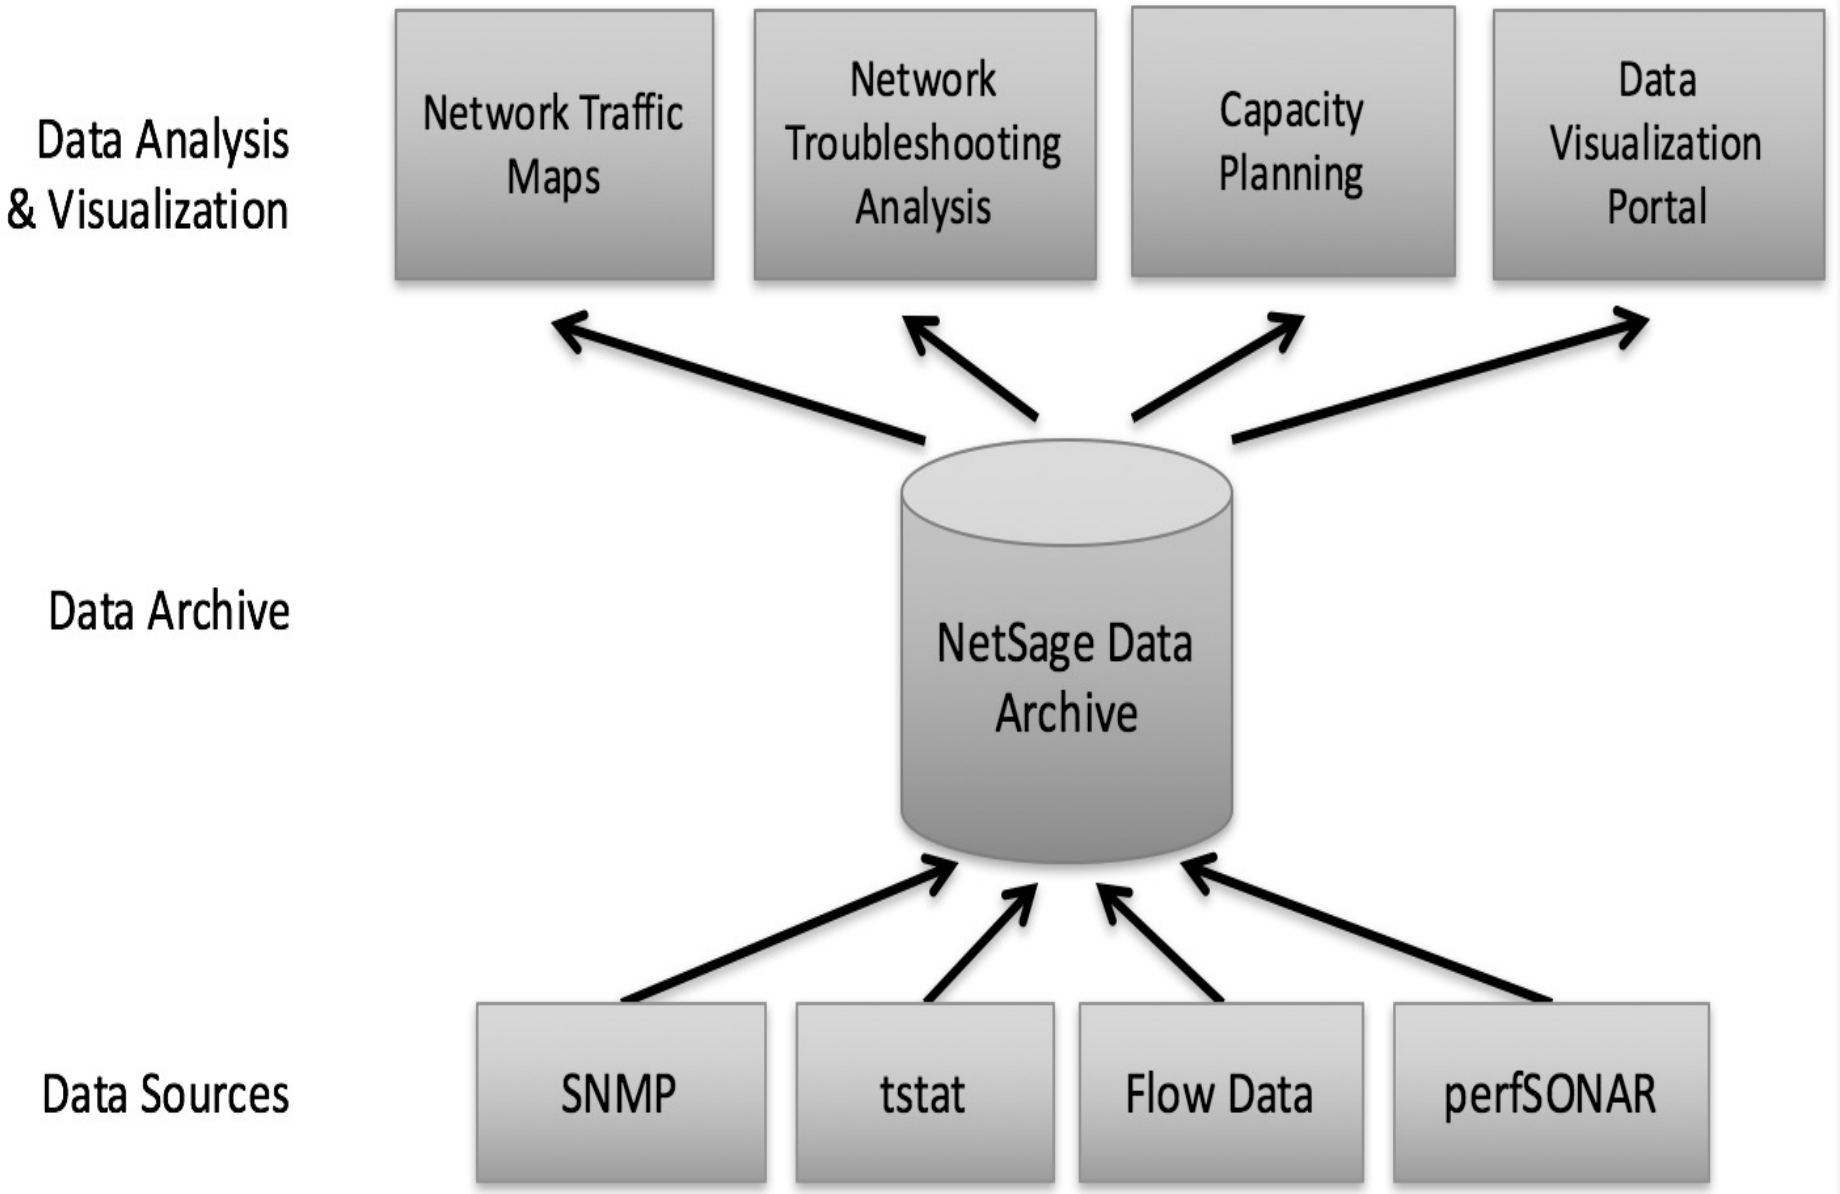

Privacy Policy

- NetSage Network Data Privacy Policy
- NetSage Project Network Data Collection Memorandum of Cooperation
  - Policy, not implementation

- IRNC PIs have commented on
- Latest draft will be posted next week

Use Cases– Year 1 & 2

- What is the max, min, average bandwidth used between links?
- Which exchange points or networks are congested?
- When and how often do they remain congested?

OPERATIONS

- What is the max, min, average bandwidth
- What is the duration and are there any periodic patterns or peak periods
- Which exchange points or networks are congested
- When and how often do the links / exchange points remain congested
- What are the top sites that use the IRNC links (average and peak)
- What are the top science projects that use the IRNC links (average and peak)

QUALIFIERS

- Location
  - Worldwide
  - Country
  - States(Regions)
  - Building
  - Specific Link
  - Institutions
- Application
  - Database/Mail…
  - UDP/TCP…
  - Audio/Video/Text…
- Disciplines
  - Physics/ Geoscience …
  - NSF Project

TIME

- Now
- Yesterday
- Last week
- Last month
- Last Year
- This year
- Specific (year, day…)
- Tagged Event: SC Conference
- Duration: time to time

IRNC Archive
Time Series Data System (TSDS)

- Common archive shared with IRNC NOC
- Open Source software on commodity hardware
- Provides well structured and high performance storage and retrieval of timeseries data
- Capable of tracking and reporting based on metadata
  - eg. viewing interface throughput from the viewpoint of a VLAN or BGP peer sessions from
- https://github.com/GlobalNOC/tsds-services

Active Measurements on Backbones

- PerfSONAR
  - Open source, community based tool for latency and throughput data
  - Hook into existing testpoints deployed on IRNC framework
  - Pull data into common archive
  - Enable common queries across IRNC sites

- Would like to include Exchange Points soon

Netsage Mesh Dashboard

Netsage Mesh - NetSage BWCTL Testing - Atlantic

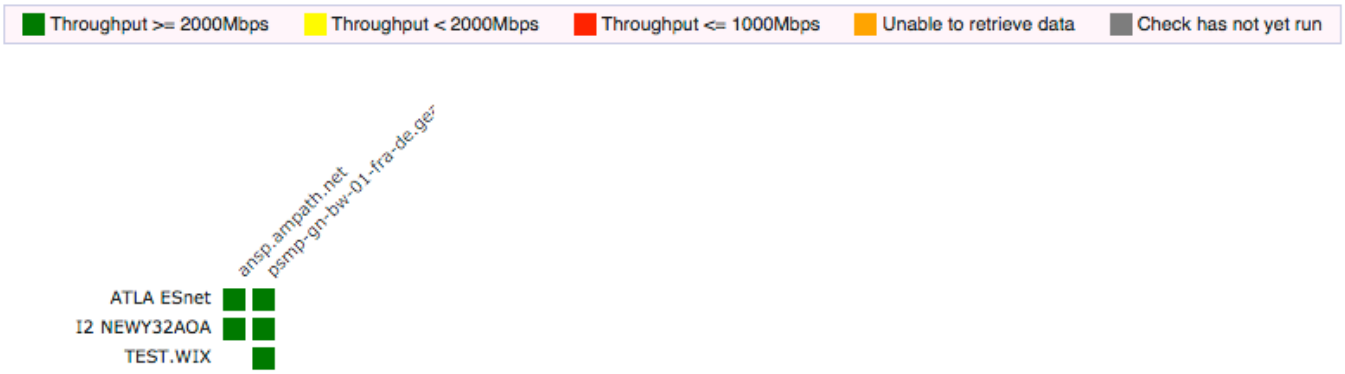

Netsage Mesh - NetSage OWAMP Test

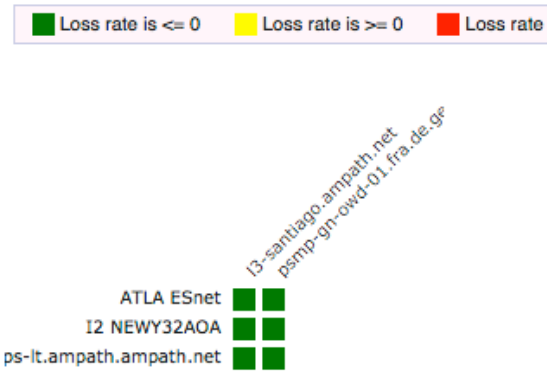

Netsage Mesh - NetSage BWCTL Testing - Pacific

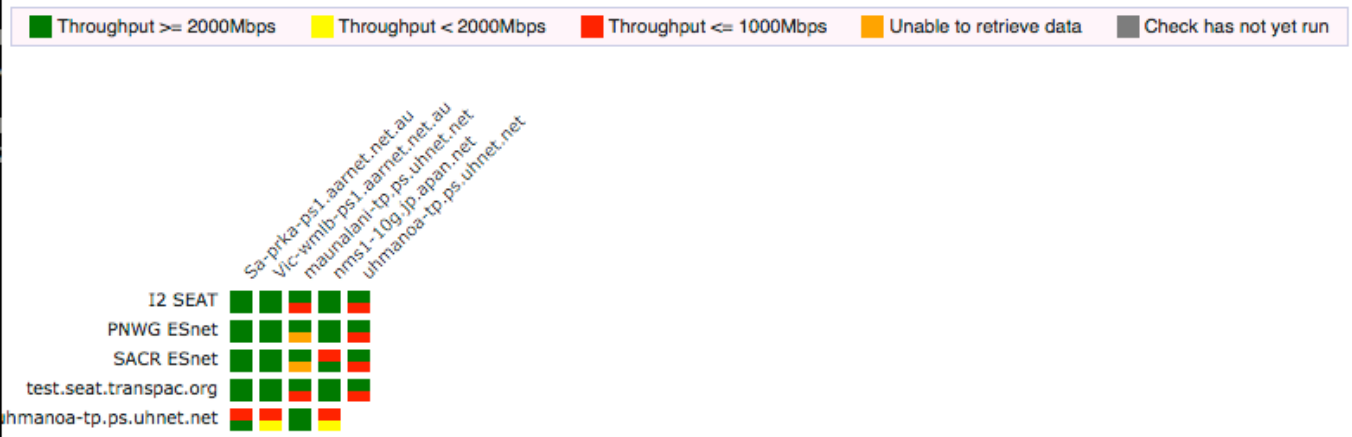

Netsage Mesh - NetSage OWAMP Test

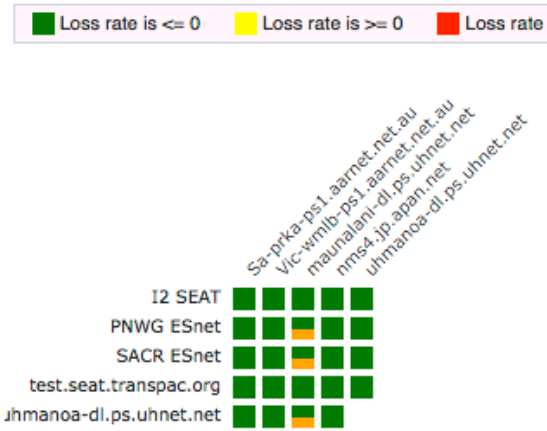

Passive Measurements on Backbones: SNMP

- All of the IRNC backbones and exchange points already collect this
- Requirement of sharing with the NOC was already in place
- Access via SNAPP tool
- Common queries across all IRNC sites

- All backbone data being collected, NOC working to extend to exchange points

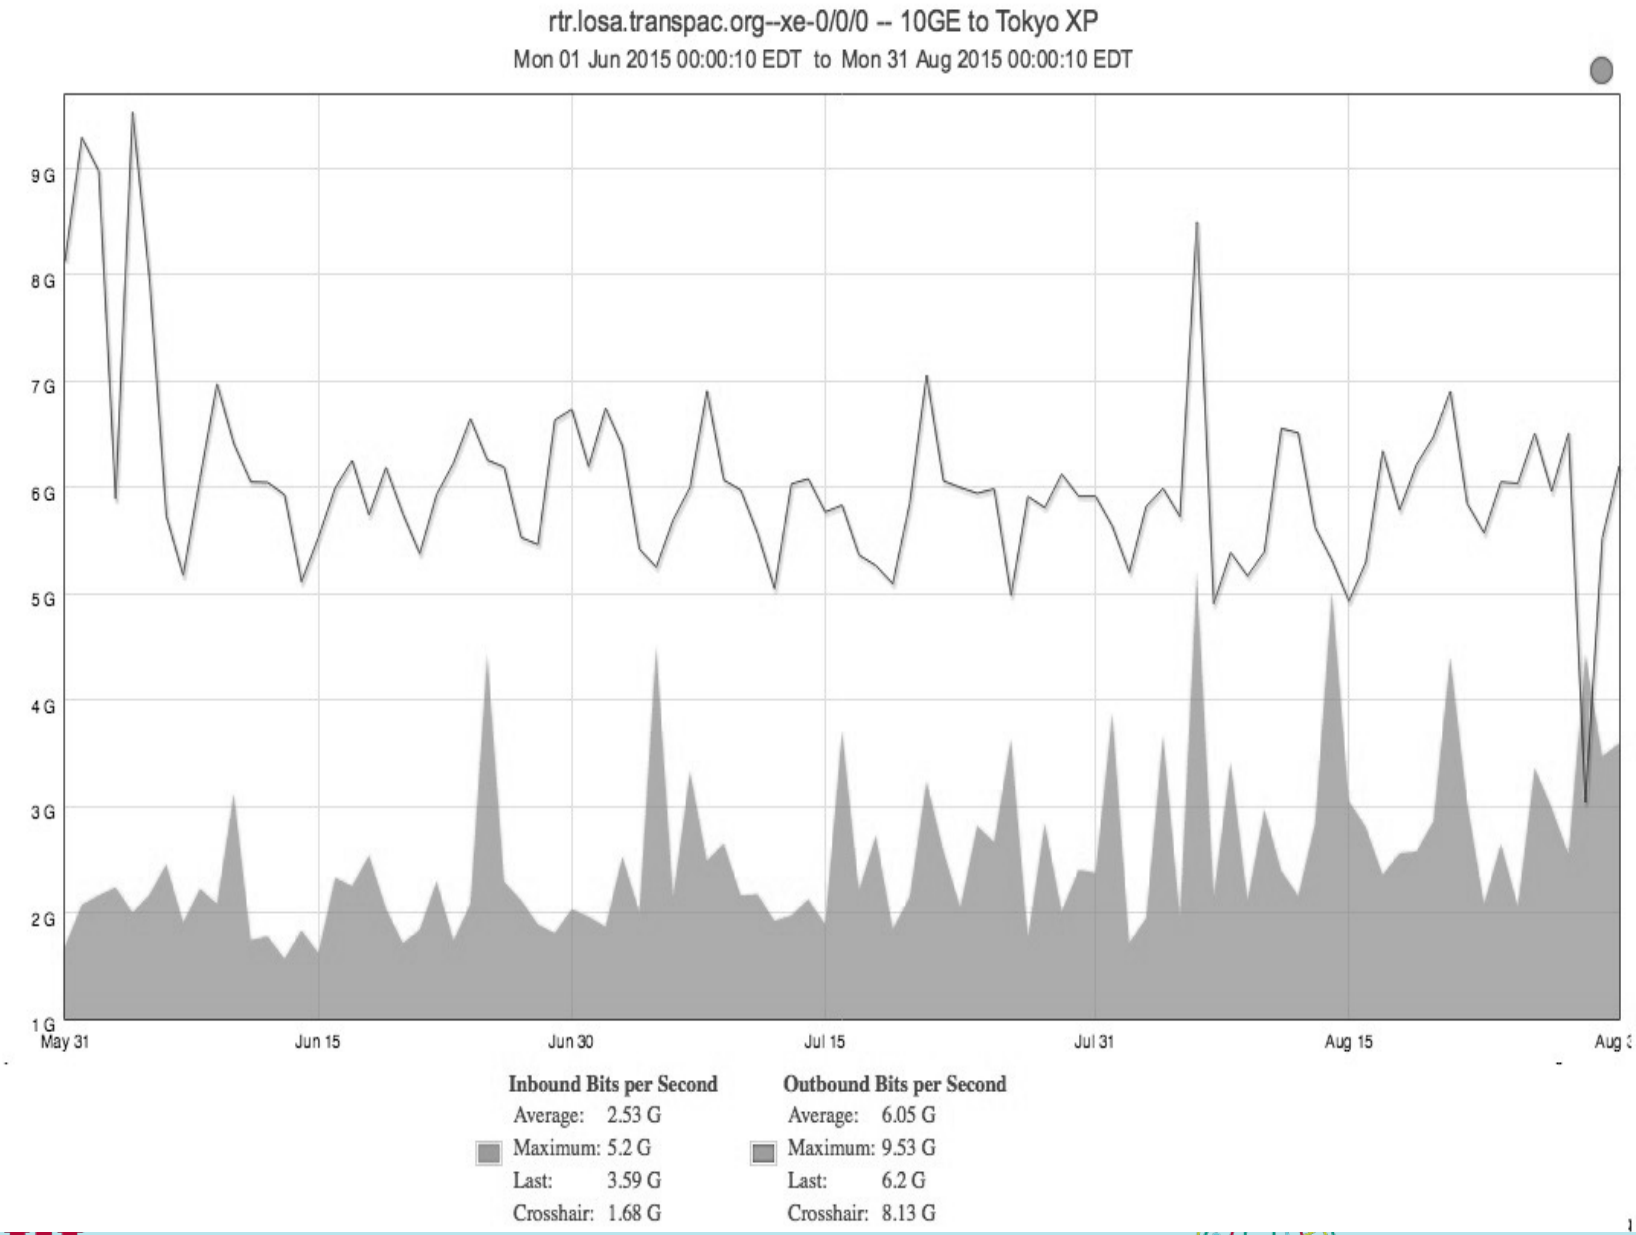
rtr.losa.transpac.org--xe-0/0/0 -- 10GE to Tokyo XP
Mon 01 Jun 2015 00:00:10 EDT to Mon 31 Aug 2015 00:00:10 EDT
9 G
8 G
7 G
6 G
5 G
4 G
3 G
2 G
1 G
May 31
Jun 15
Jun 30
Jul 15
Jul 31
Aug 15
Inbound Bits per Second
Average: 2.53 G
Maximum: 5.2 G
Last: 3.59 G
Crosshair: 1.68 G
Outbound Bits per Second
Average: 6.05 G
Maximum: 9.53 G
Last: 6.2 G
Crosshair: 8.13 G

Passive Measurements
De-Identified Flow Data

- TransPAC was first gyuinea pig
  - Was already collecting de-identified flow data
- Adapted archive and queries to include this format
- Adding analysis tools to portal

Evaluation of Deep Packet Inspection Tools Bro, tstat

- Bro couldn't scale for the bandwidth we needed
  - Plus path asymmetry issues
- Tstat, part of the EU 'mplane' (Measurement Plane) project
  - http://tstat.polito.it/
  - Able to do analysis of 10G TCP flows using <10% of a single core
  - Includes ability to de-identify IP using Crypto-Pan prefix-preserving IP anonymizer
- Use tstat for TCP retransmission analysis and basic flow data

Current Status Flow Data

- TransPAC:
  - sFLOW (both ends) and Tstat
- Ampath
  - sFlow expected before SC
- PIREN
  - In discussions for best tool
- ACE
  - SFLOW and tstat
- Exchange points later in Year 2/3

Ask NetSage...

Current IRNC links with bandwidth of links and exchange points displayed on a map.

What was the min, max, average
in Bandwith use
across the IRNC Network
time frame From: 01/01/2016 at: 00:00 to: 09/13/2016 at: 02:00

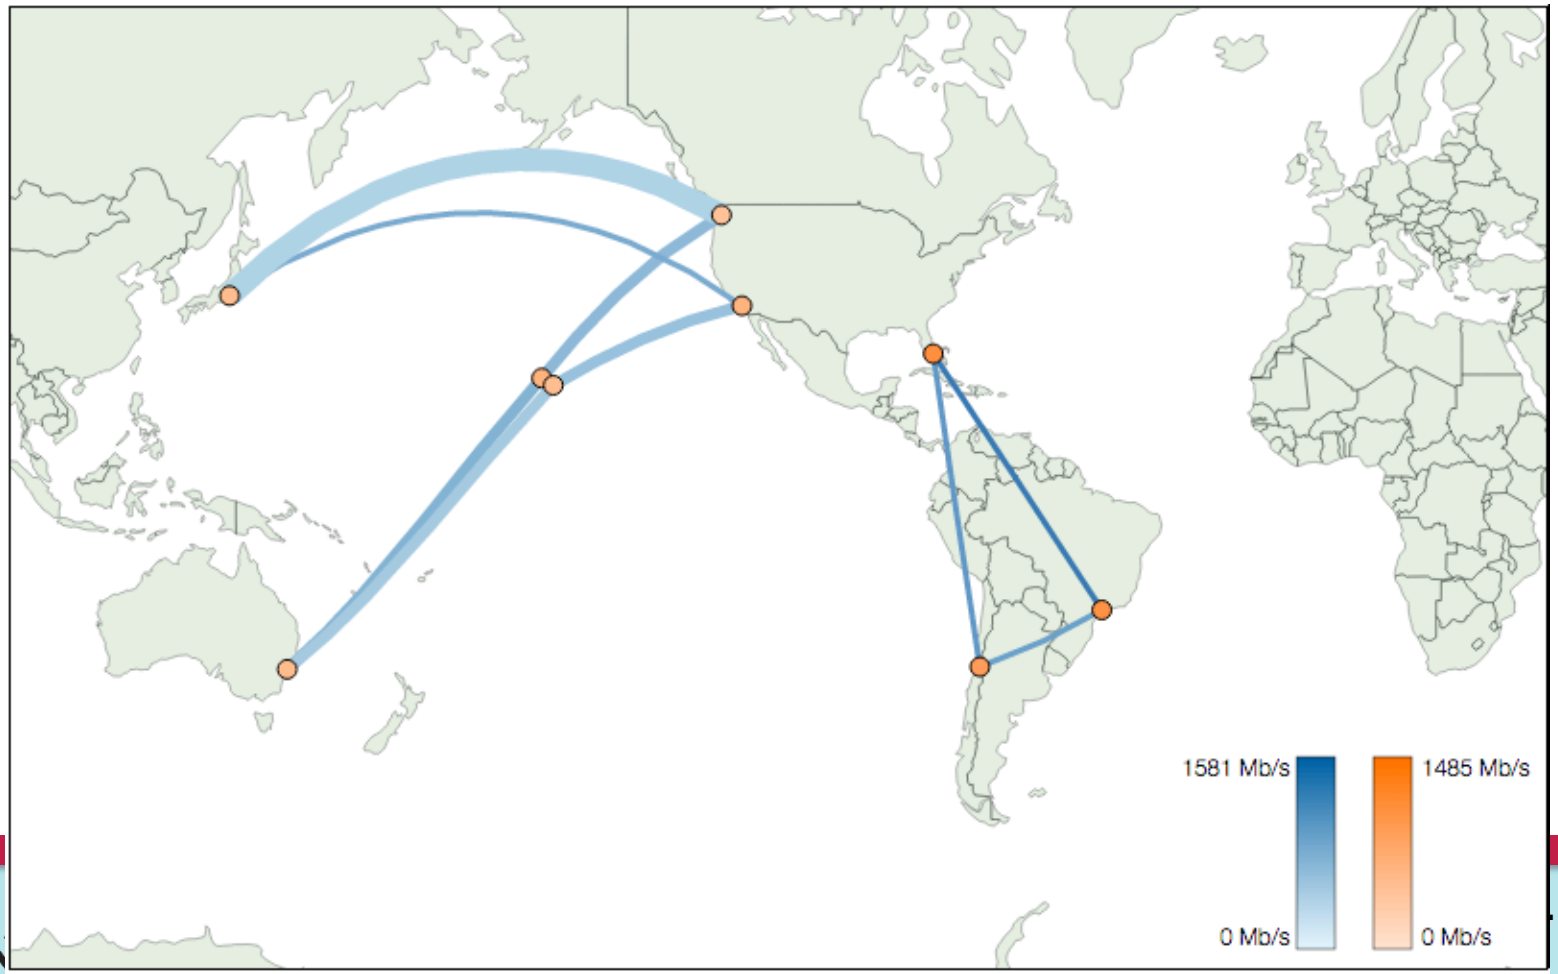

Ask NetSage...

What was the min, max, average in Bandwith use across the IRNC Network time frame From: 01/01/2016 at: 00:00 to: 09/13/2016 at: 02:00

Display of incoming and outgoing bandwidth & total data transmitted **between links**.

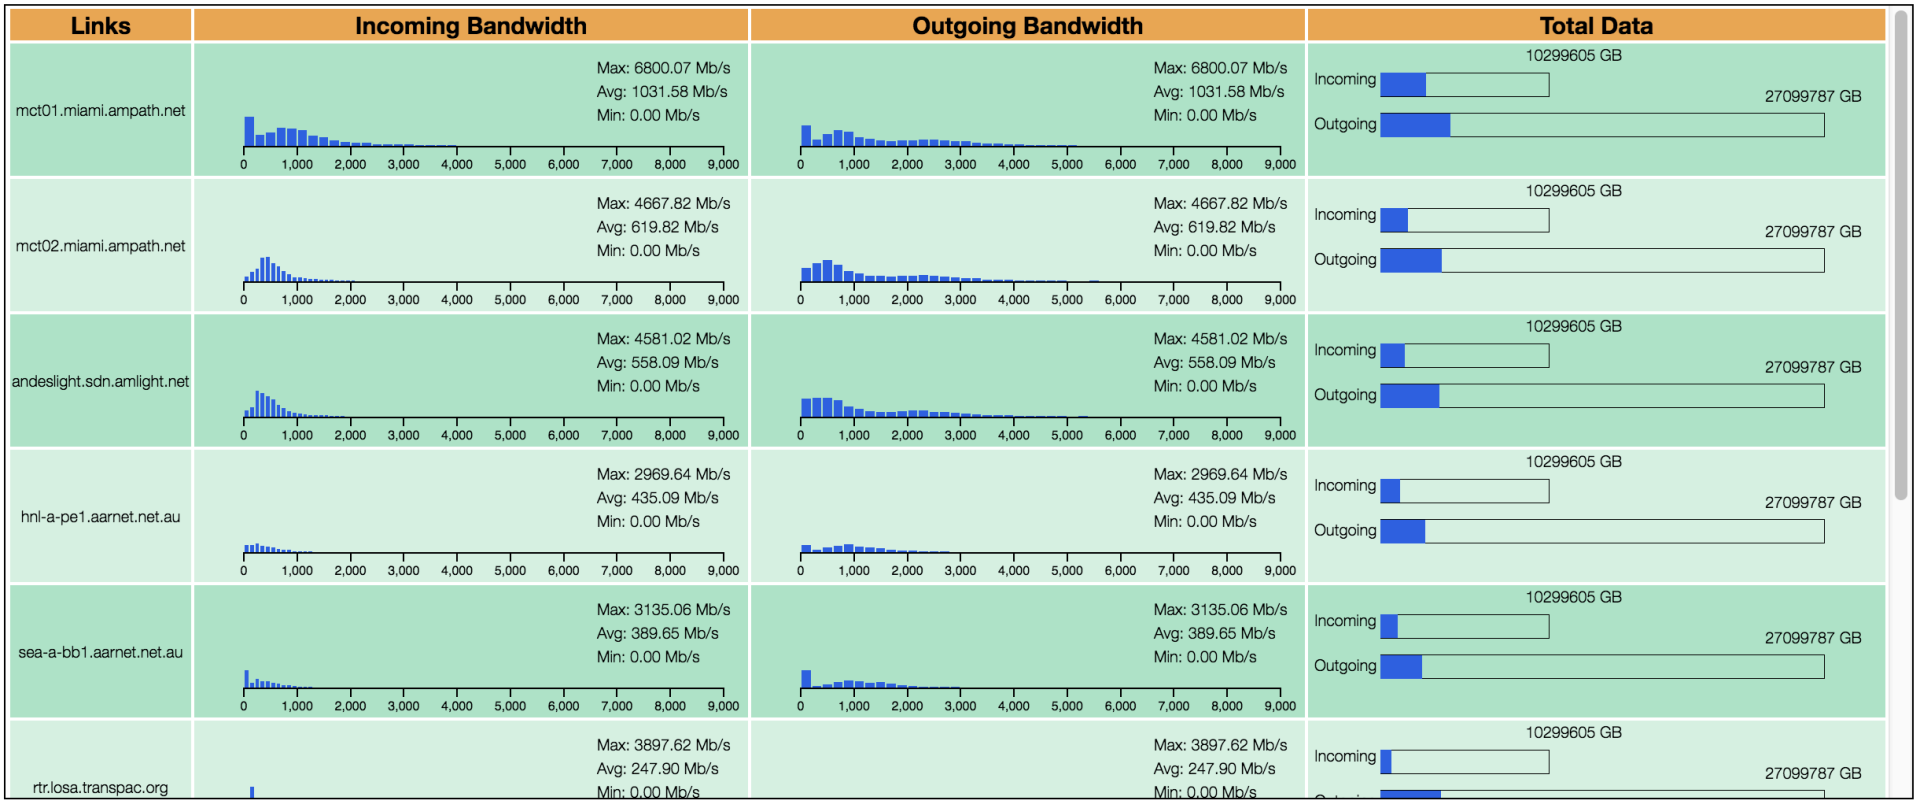

Distribution of bandwidth use for each link

Total transmitted data for each link relative to all the links

INDIANA UNIVERSITY

INTERNATIONAL NETWORKS
At Indiana University

Ask NetSage...

What was the min, max, average in Bandwith use across the IRNC Network time frame From: 01/01/2016 at: 00:00 to: 09/13/2016 at: 02:00

Display of incoming and outgoing bandwidth & total data transmitted in **exchange points**.

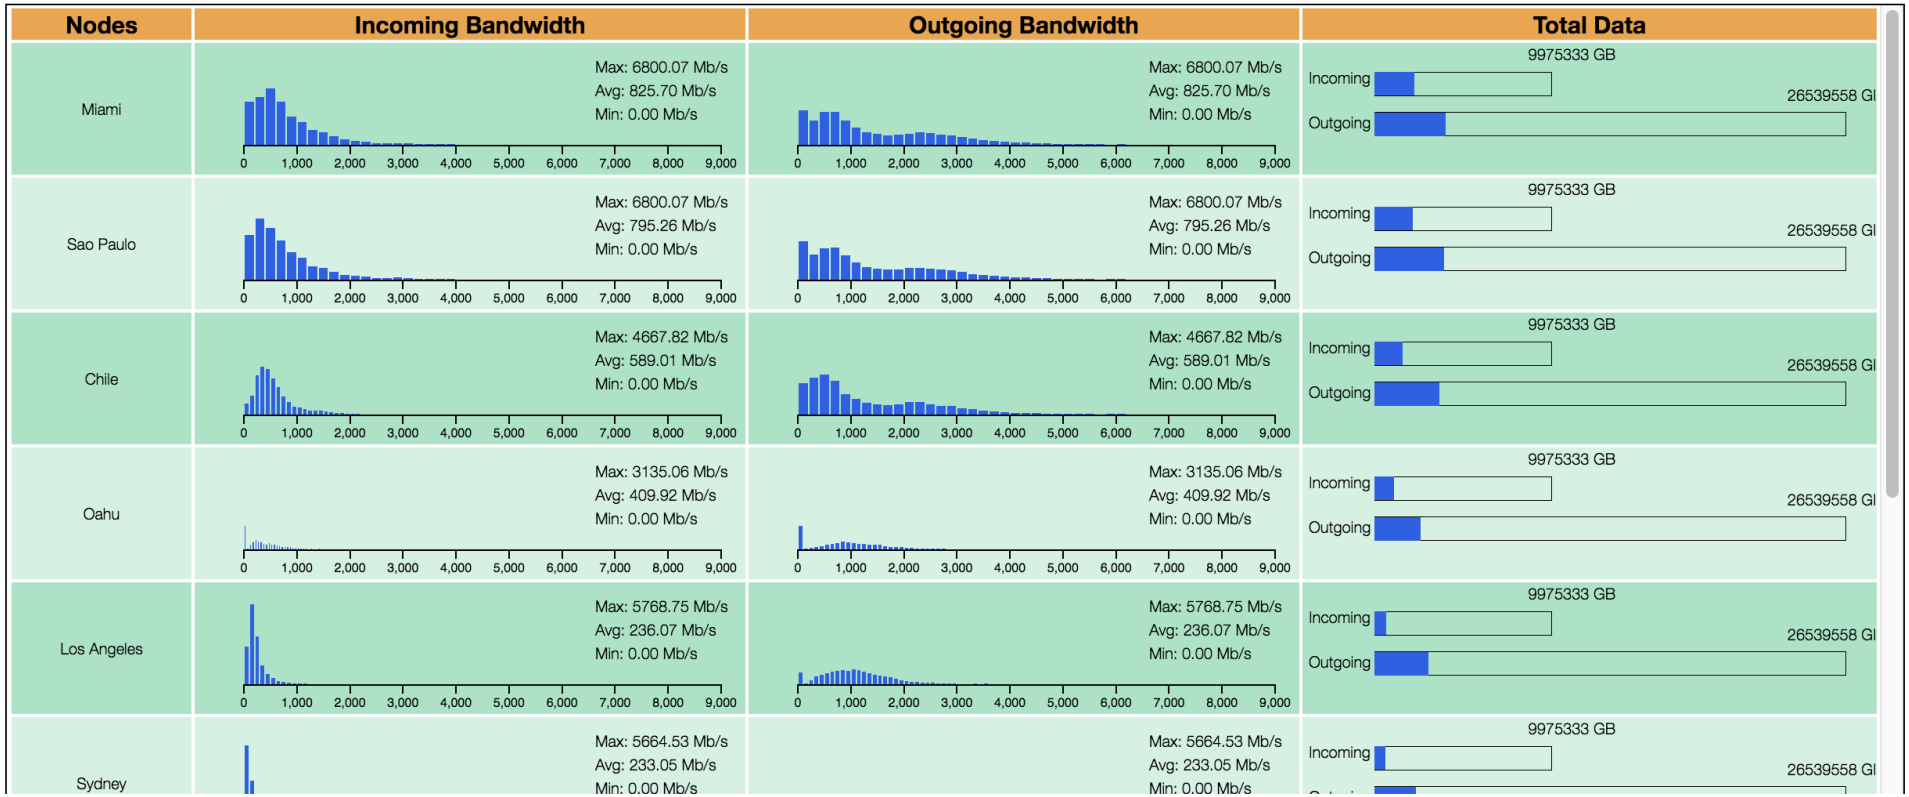

Distribution of bandwidth use for each exchange point

Total transmitted data for each link relative to all the exchange points

INDIANA UNIVERSITY

INTERNATIONAL NETWORKS
At Indiana University

Ask NetSage...

What was the min, max, average
in Bandwith use
across the IRNC Network
time frame From: 01/01/2016 at: 00:00 to: 09/13/2016 at: 02:00

Heatmaps showing periodic patterns in bandwidth use across the IRNC networks from 8/1/16-9/1/16

What is the duration and are there any periodic patterns or peak periods in Bandwith use across the IRNC Network time frame: 08/01/2016 10:00:00 UTC , 09/01/2016 09:00:00 UTC

Input: mct02.miami.ampath.net

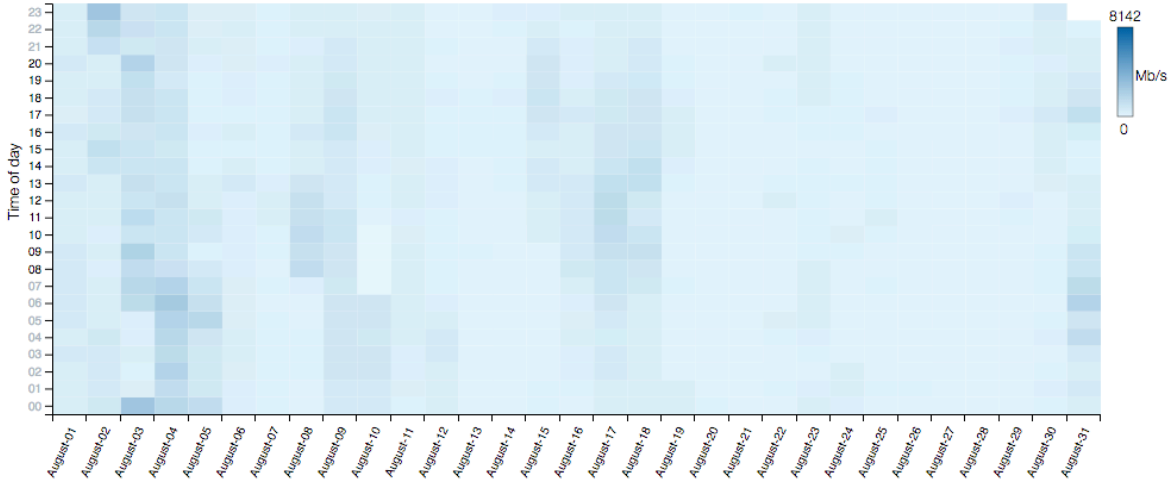

Aggregated weekly bandwidth use use

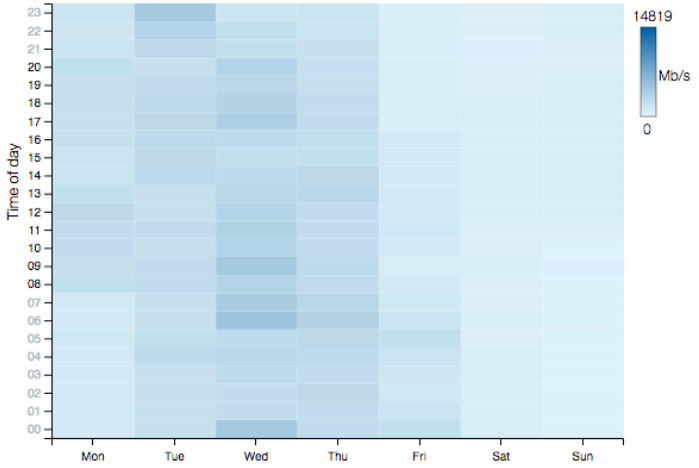

Output: mct02.miami.ampath.net

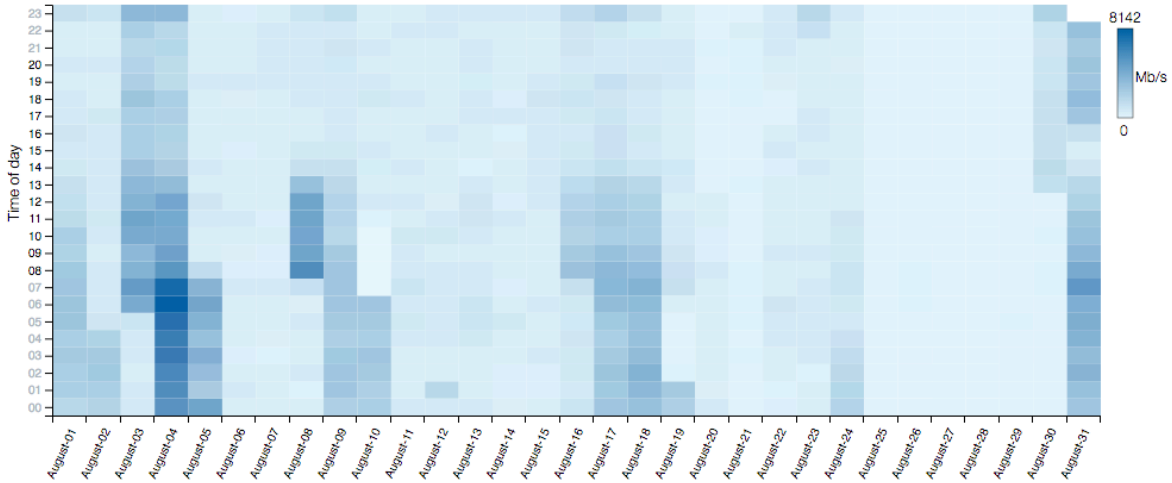

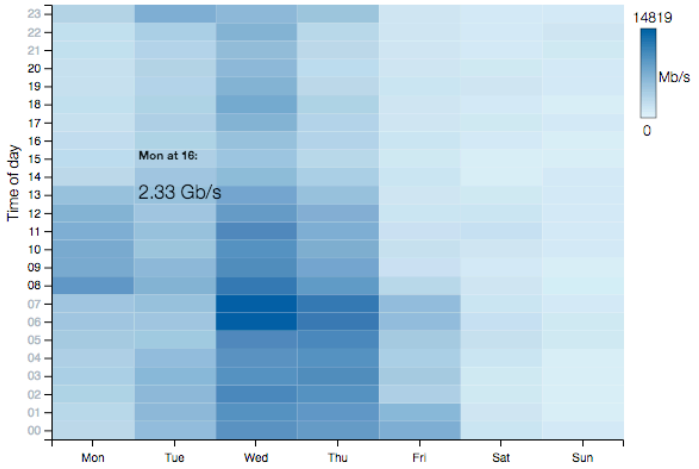

Next steps

- Expanding the data collection
  - What data will you be collecting?
- Analysis
  - What questions are you trying to answer?

Questions/Comments?

- **NetSage Website:**
  - http://www.netsage.science

- **Questions? Contact**

Jennifer Schopf – jmschopf@indiana.edu