

Global Science Engagement: What's Working and Paths Forward

- **Eli Dart, ESnet**
 - Engagement With Scientific Facilities
- **David Wilde, AARNet**
 - An NREN is not an island: solving researcher problems through partnership
- **Lloyd Ssentongo, RENU**
 - A Successful collaboration between NIH ICER Uganda and the National NREN RENU
- **Jennifer Schopf, Indiana University**
 - Science Engagement Internationally

International Science Engagement by IN@IU

Jennifer M. Schopf, PhD

Director, International Networks

August 8, 2017

“Why is understanding performance hard?”

- **Why do scientists find it hard to move data effectively?**
- **Why is it hard to find the right network person to talk to about problems?**

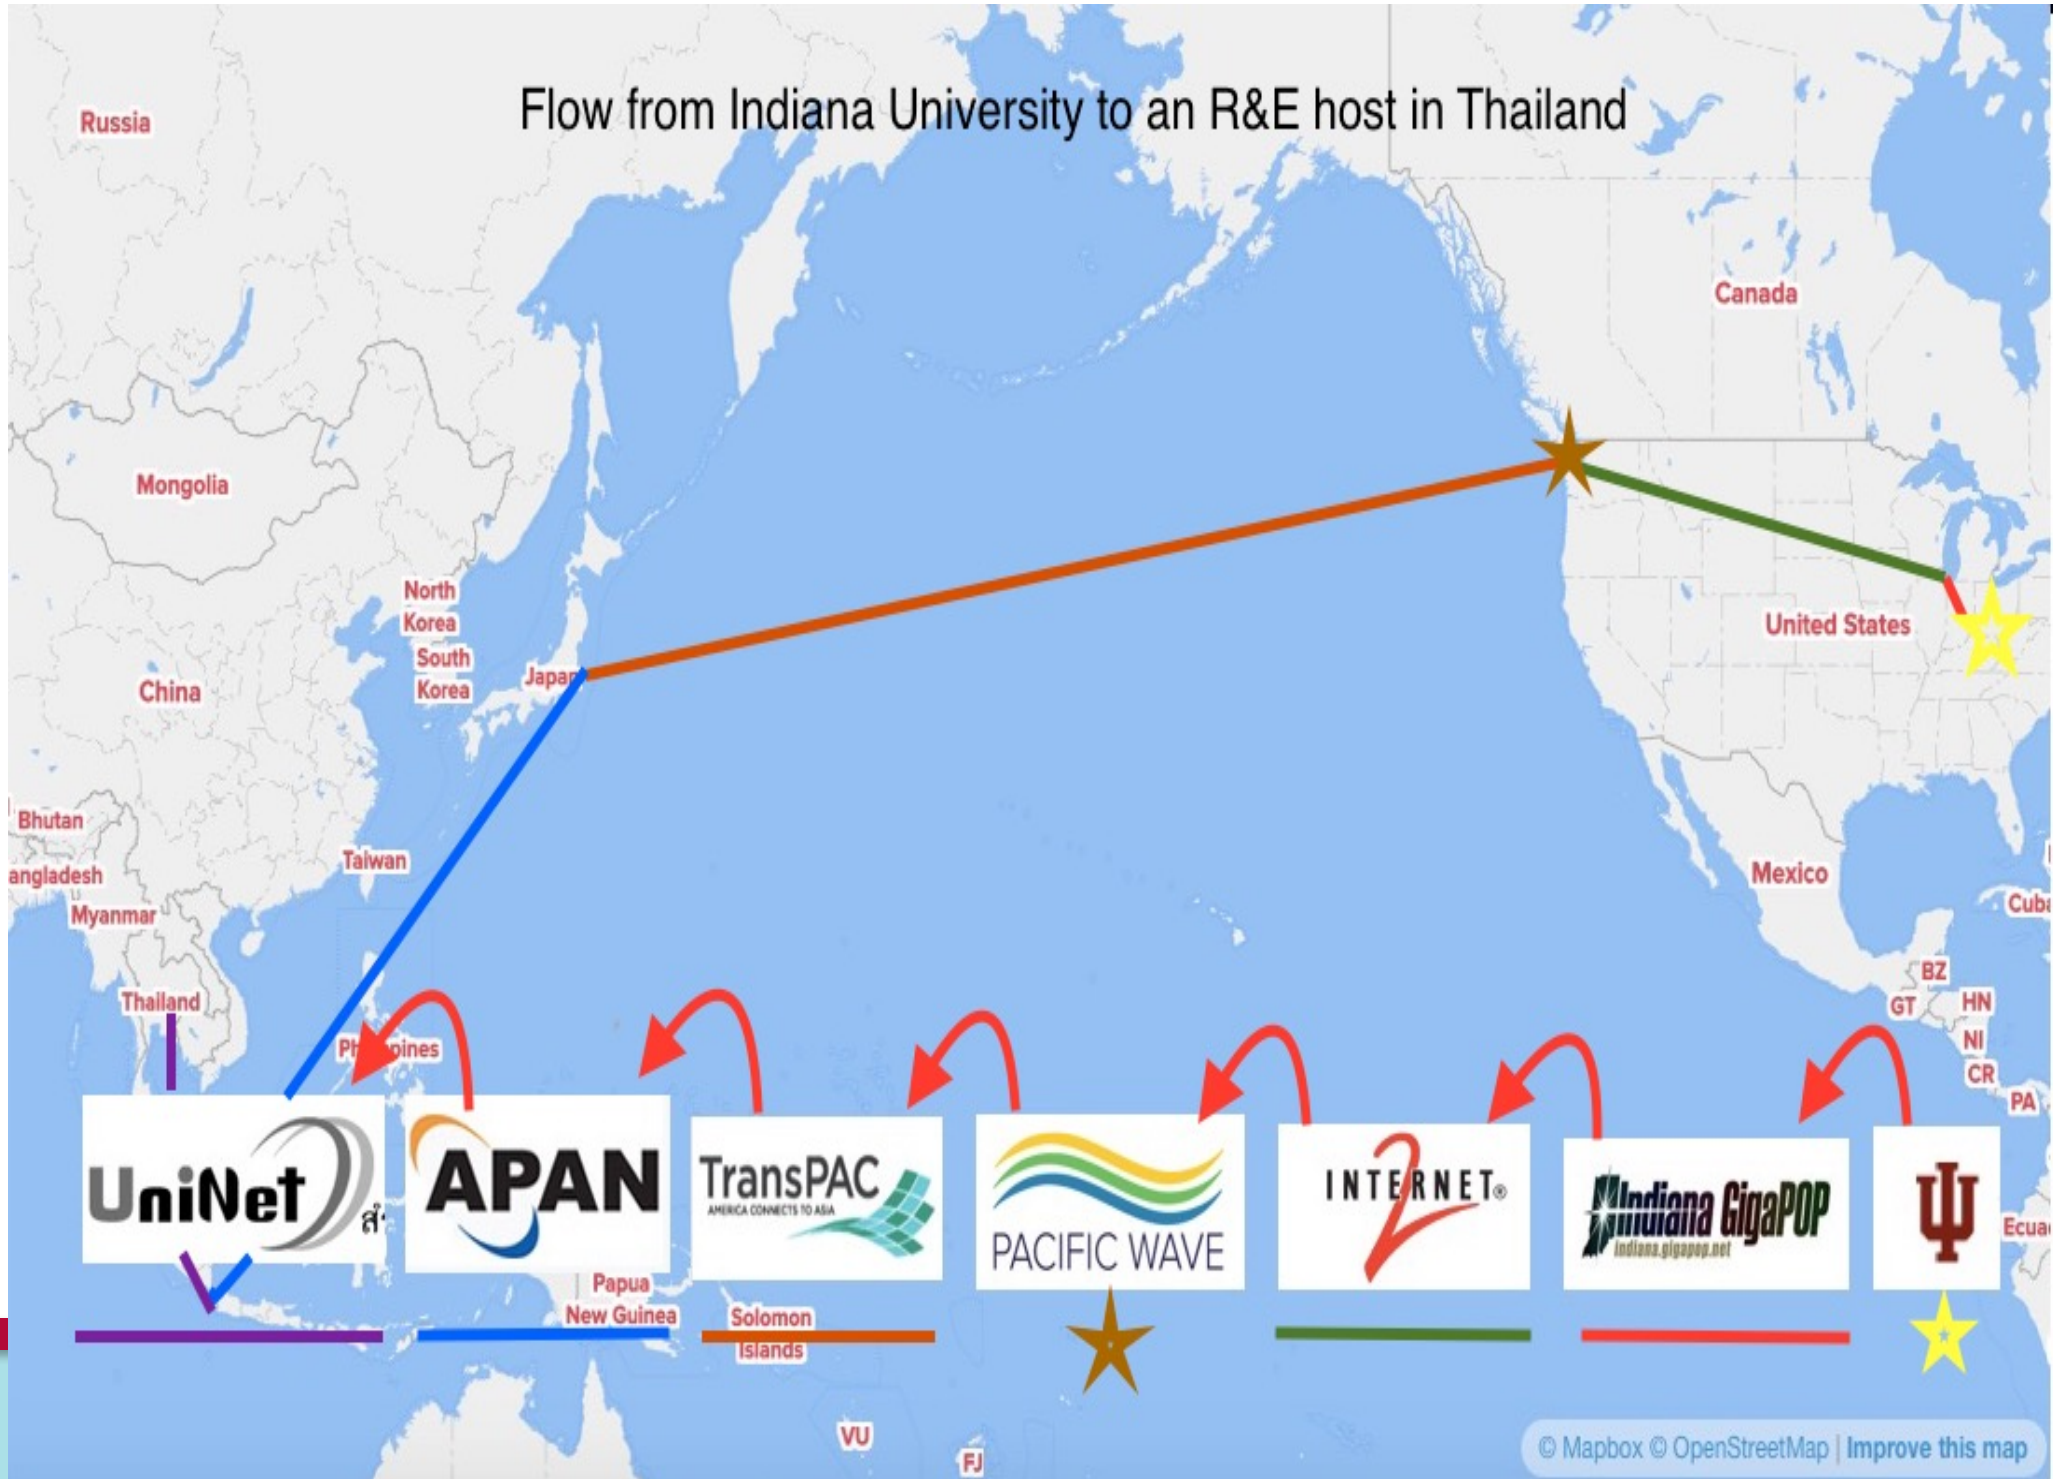
Why is understanding data movement hard?

- **Lots of pieces - Host system through networks to host system**
- **No one controls all the pieces**
- **Unknown expectations for what performance should be**
 - **Theoretical vs pragmatic**
- **Asymmetric paths**
- **Soft failures are hard to find**
- **Many, many points of coordination**

Flow from University of Utah to University of Cape Town



Flow from Indiana University to an R&E host in Thailand



Two Types of Engagement


- **Case Studies**
 - Walk through your science
 - Understand needs and planning
 - Think of this as regular maintenance, oil change, or planning to buy a new car
- **“Roadside” Assistance**
 - This file transfer worked last week!
 - Think of this like a flat tire, crash repair
 - At IU this is done by both IN@IU and the IU Performance Engagement Team (PET)

Case Studies

- **Jason Zurawski, ESNNet**
- <https://fasterdata.es.net/science-dmz/science-and-network-requirements-review/>
- **We did this in March at CENIC meeting with Steve Diggs, Scripps Oceanography**

Anatomy of the Case Study

- There are two primary components to this exercise
 - Narrative
 - Data Estimation
- ESNet has developed a ‘template’
 - Contains helper text to guide what is wanted
- Items that make sense for ESnet, may not make sense for your institution. Modify as needed
- <https://fasterdata.es.net/science-dmz/science-and-network-requirements-review/>

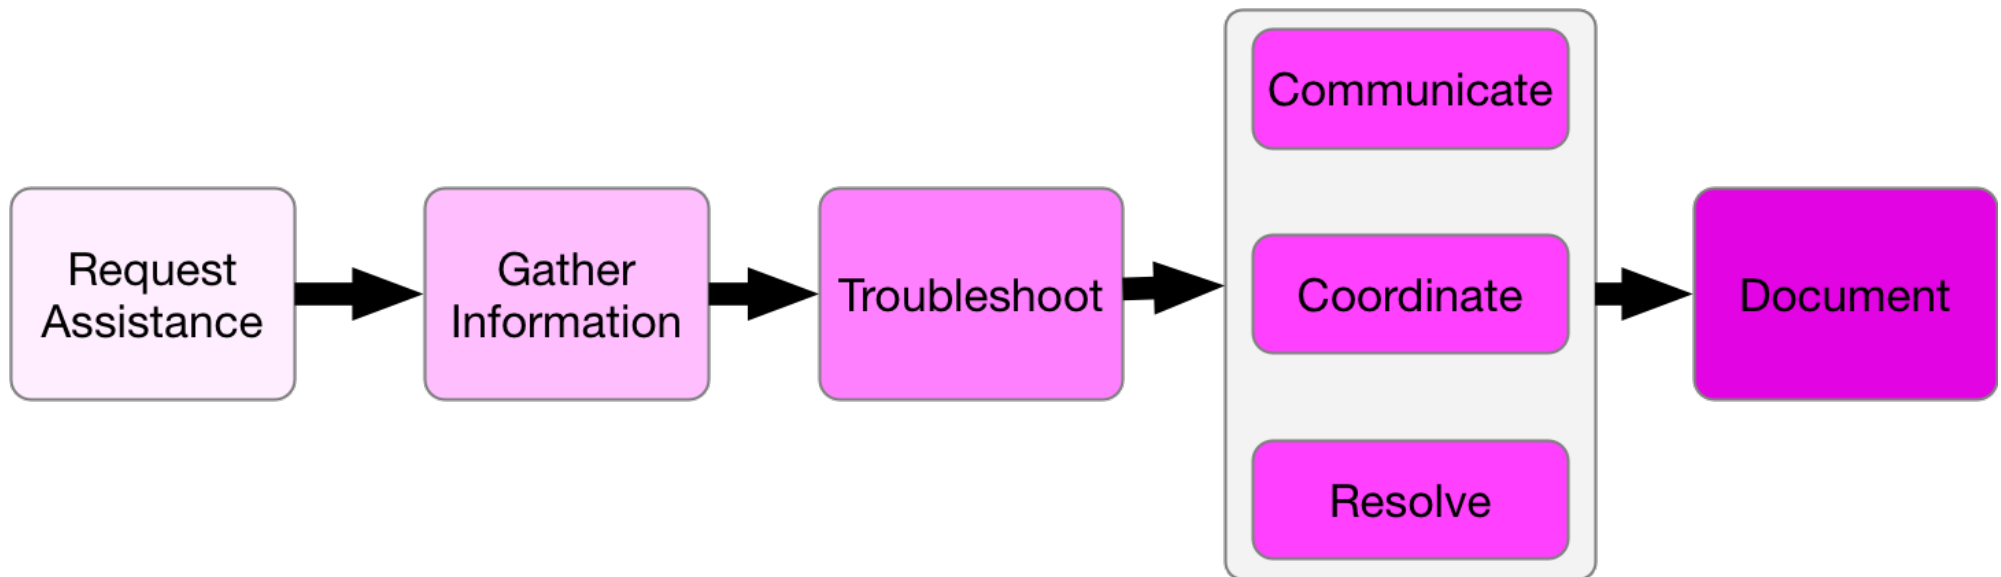


The Narrative (Section Overview)

- Title & Author Information
- Background
- Collaborators
- Instruments and Facilities
- Process of Science
- Remote Science Activities
- Software Infrastructure
- Network and Data Architecture
- Cloud Services
- Outstanding Issues

IU Performance Engagement Team

- **Help researchers troubleshoot network performance by providing centralized support along the end-to-end path**

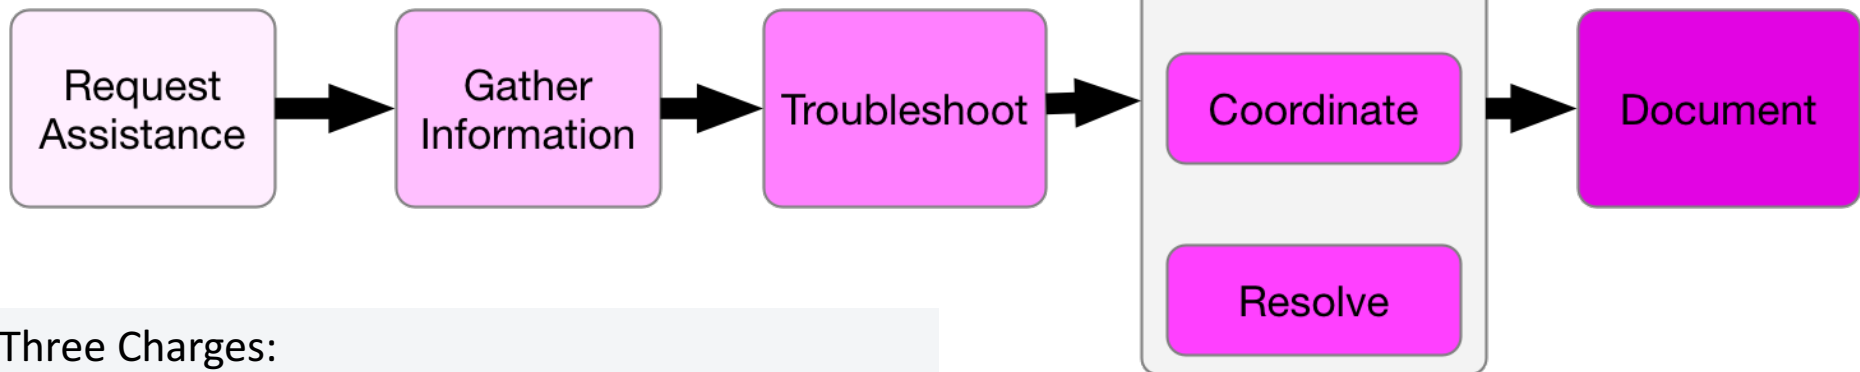


IRNC Performance Engagement Team

The IRNC PET exists to help researchers troubleshoot network performance by providing centralized support along the end-to-end path.

Network Paths have become more complex:

- Layer2 infrastructure obscures the network path
- Heightened security requirements mean there's less publicly-available metrics
- Network firewall usage on the rise
- Automated data transfer requires 24x7x265 support



Three Charges:


- 1) Drive quick resolution of inter-domain performance issues
- 2) Build a common troubleshooting playbook
- 3) Evolve personar as a tool for performance incident management

Anyone can submit a problem!

<https://irncnoc.globalnoc.iu.edu/performance@globalnoc.iu.edu>

Debugging the TransPAC Link

- **perfSONAR's owamp tool detected packet loss on the path from Seattle to Japan**
 - **No other monitoring detected this**
- **Opened a ticket with PET**
 - **Built a map (traceroute)**
 - **Pulled in collaborators**
 - **PNWGP, WIDE, APAN-JP, Internet2, and others**
 - **Got extra data from other PS nodes**
 - **Ran some additional tests**



And discovered

Packets were getting dropped by a Brocade in Seattle

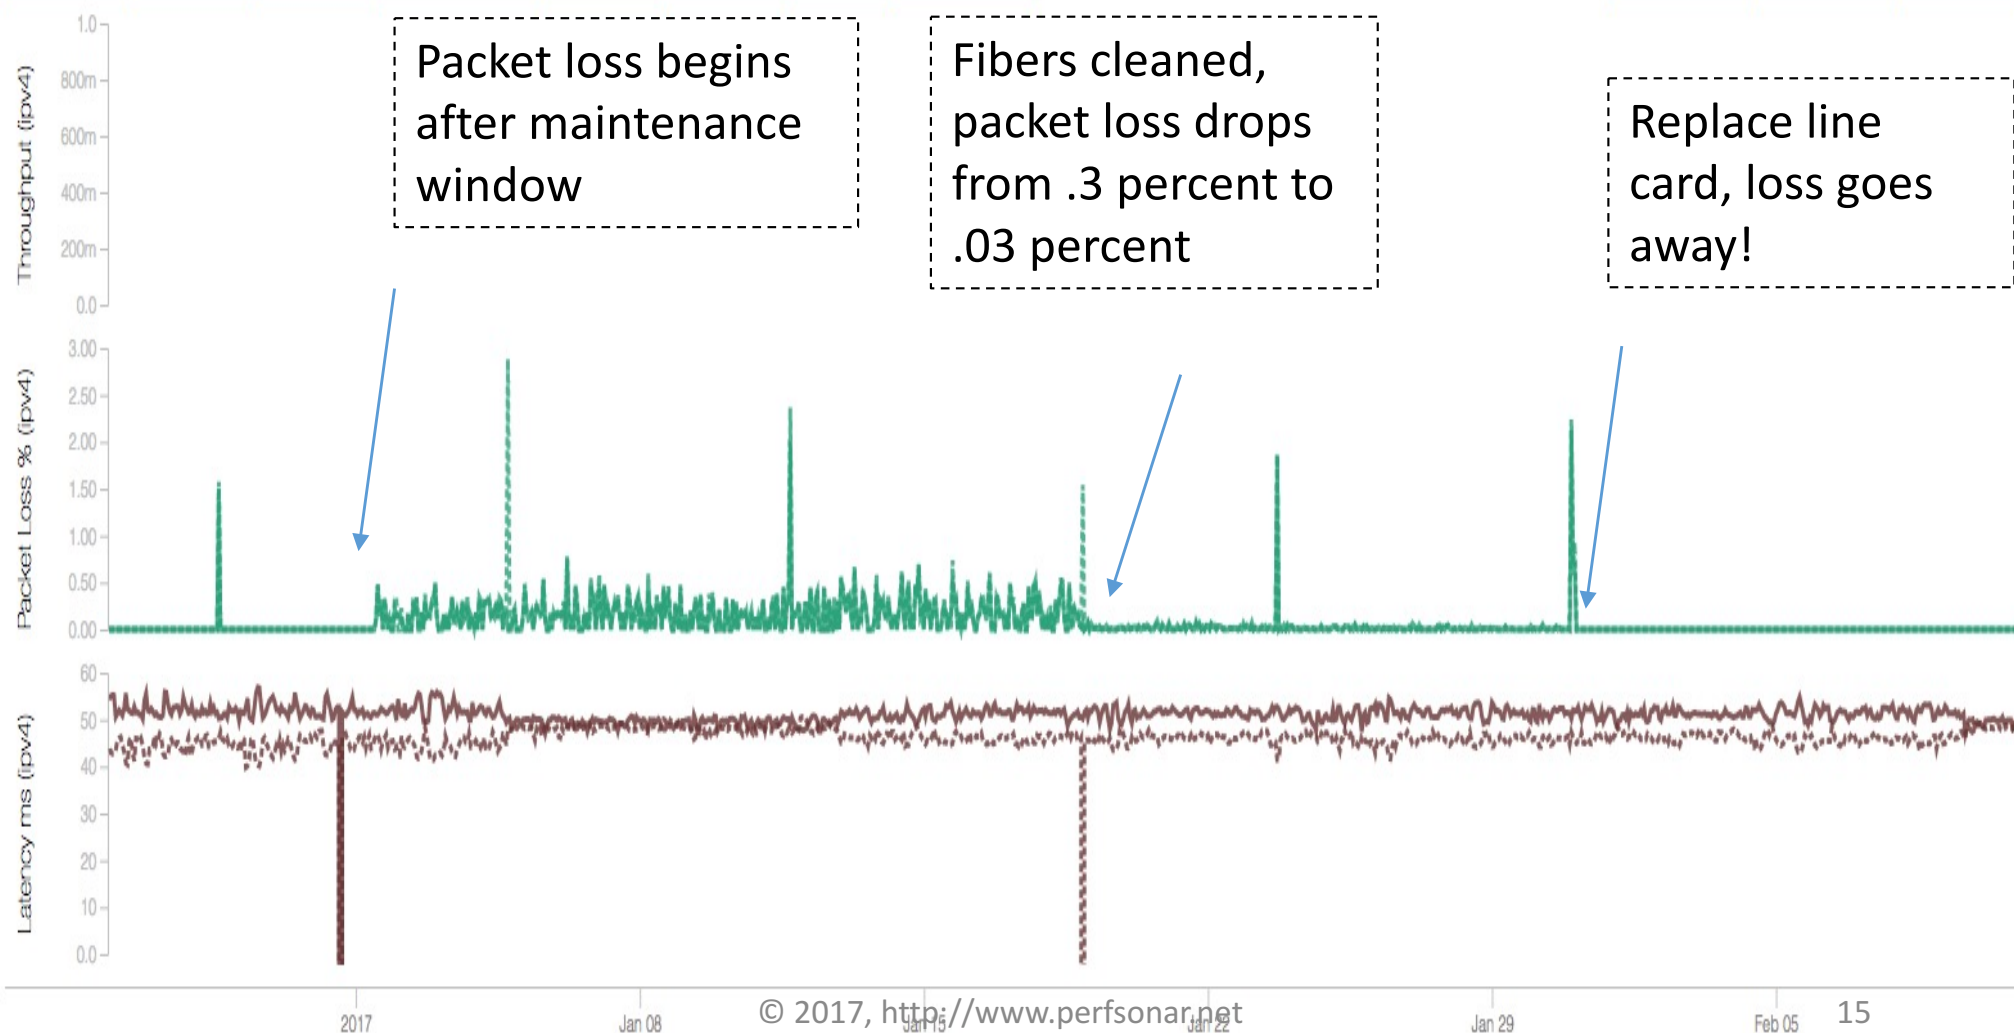
- **Cleaned fibers: Problem got better, but not gone**
- **Replaced a defective line card**

Source
test.seat.transpac.org
192.203.115.2
[Host info](#) ▾

Destination
nms4.jp.apan.net
203.181.249.197
[Host info](#) ▾

Report range
← 1 month ▾ →
Sun 12/25/2016 to Sat 02/11/2017
21:18:02 (GMT-8) 02:11:22 (GMT-8)

Tput (TCP) Tput (UDP) Loss (UDP) Loss (owamp) Loss (ping) Retrans Latency Latency (ping) Forward Reverse Failures



GWU-Netherlands Performance Issue

- Context: Professor at George Washington University needs to move >40T + incremental updates from site in the Netherlands but is not getting expected performance

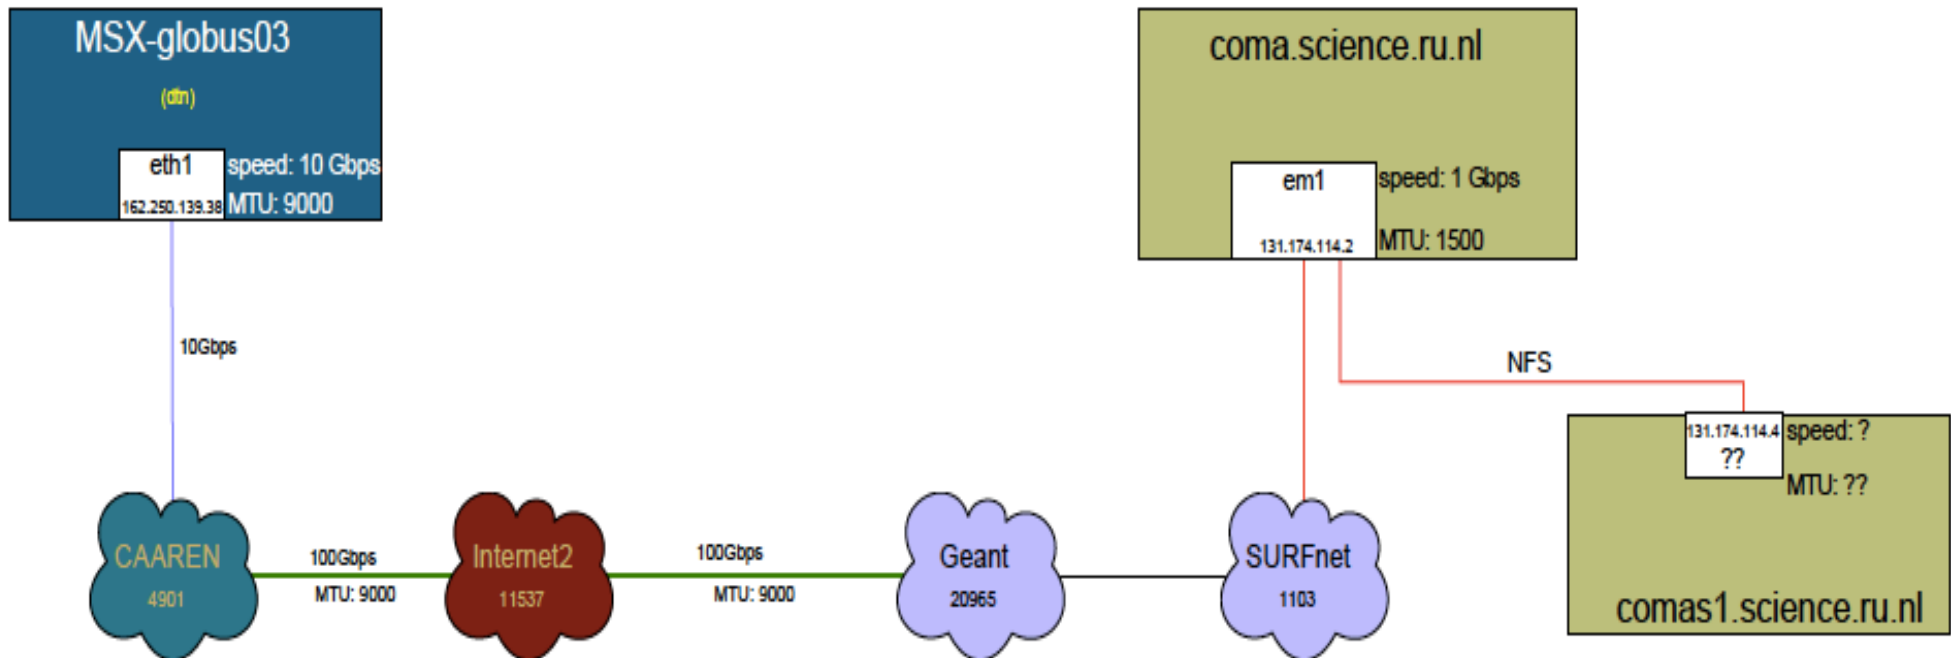
Problem:

- LOFAR end host is not optimized for big data transfers
 - Only a 1G interface
 - Shared system, other users running jobs
 - Data resides on a NFS mount only reachable over the same 1G interface
- GWU CIO reaches out to ASTRON to try to improve performance but is unsuccessful

GWU-Netherlands Performance Issue

- **Context: Professor at George Washington University needs to move >40T + incremental updates from site in the Netherlands but is not getting expected performance**
- **Problem:**
 - **LOFAR end host is not optimized for big data transfers**
 - **Only a 1G interface**
 - **Shared system, other users running jobs**
 - **Data resides on a NFS mount only reachable over the same 1G interface**
- **GWU CIO reaches out to ASTRON to try to improve performance but is unsuccessful**

GW LOFAR TRANSFER TOPOLOGY



```
/vol/astro5/lofar/msss/raw_data/
```

```
/vol/astro5/lofar/msss/raw_data/hba-201*/set*/H*/sky.model
/vol/astro5/lofar/msss/raw_data/hba-201*/set*/H*/BAND*/*.MS
```

```
/vol/astro5/lofar/msss/raw_data/hba-20130615/H*/sky.model
/vol/astro5/lofar/msss/raw_data/hba-20130615/H*/BAND*/*.MS
```

Role of IN@IU Science Engagement Team

- Worked with GWU to understand their performance issues
- Contacted SURFnet (Netherlands) to understand local connectivity and their relationship with ASTRON
- Works with SURFnet and ASTRON to isolate issues:
 - Data is being pulled from multiple locations in the Netherlands
 - ASTRON team only has control over performance from the LOFAR LTA databases
 - Some data has been migrated to tapes at university locations

Solution:

- IN@IU and SURFnet are working with ASTRON and Radboud University to establish Science DMZs for LOFAR data at multiple locations in the Netherlands

Role of IN@IU Science Engagement Team

- **Worked with GWU to understand their performance issues**
 - **Contacted SURFnet (Netherlands) to understand local connectivity and their relationship with ASTRON**
- **Worked with SURFnet and ASTRON to isolate issues:**
 - **Data is being pulled from multiple locations in the Netherlands**
 - **ASTRON team only has control over performance from the LOFAR LTA databases**
 - **Some data has been migrated to tapes at university locations**
- **Solution:**
- **IN@IU and SURFnet are working with ASTRON and Radboud University to establish Science DMZs for LOFAR data at multiple locations in the Netherlands**

Takeaways

- **Measurement, monitoring, and engagement are linked**
- **The people problem is as hard or harder than the technical ones**
- **More information: Jennifer Schopf**
jmschopf@iu.edu

Acknowledgements

- **IN@IU is funded by**
 - **NSF awards #0962968, #0962973, #1450904, #1540933, #1638863**
- **The TransPAC PacificWave 100gb/s network fabric is provided by Pacific Northwest GigaPop**
- **More information: jmschopf@iu.edu**