

How Regional Partnerships with National
Performance Engineering and Outreach
Initiatives are Enabling Science

National Performance Engineering and Outreach Initiatives: CC* 18-508

- “NPEO program area will establish a national resource of expertise, best practices, training, practical experience, data, and applied research and development”
- Goal: order-of-magnitude improvements in data transfer performance
- Two awards made:
 - #1826967: Towards a National Research Platform
 - Larry Smarr, Philip Papadopoulos, Frank Wuerthwein, Tajana Rosing, Ilkay Altintas
 - #1826994: Engagement and Performance Operations Center (ReSEC)
 - Jennifer Schopf, Jason Zurawski, Dave Jent

Solicitation Specified Activities

- Centralized point of expertise, advice, and engagement for scientists
- Investigation of and solution engineering for end-to-end performance
- Disseminating best current practices for end-to-end data transfers
- Investigating, evaluating, testing data movement systems and tools
- Providing analysis and objective recommendations
- Providing training on network performance analysis and troubleshooting
- Engaging with CI engineering expertise (eg Campus Champions, CaRCC)
- Leverage PerfSonar, passive measurement, and other performance testing
- Analyze network performance and data transfer metrics
- Identifying potential bottlenecks in data transfers for scientific collaborations

Expected collaborations with related activities and groups in R&E networking community

- Campuses
- National Research and Education Networks (NRENs)
- State and regional R&E networks
- Also “collaborate with operational entities to assist in optimizing scientific data transfers”

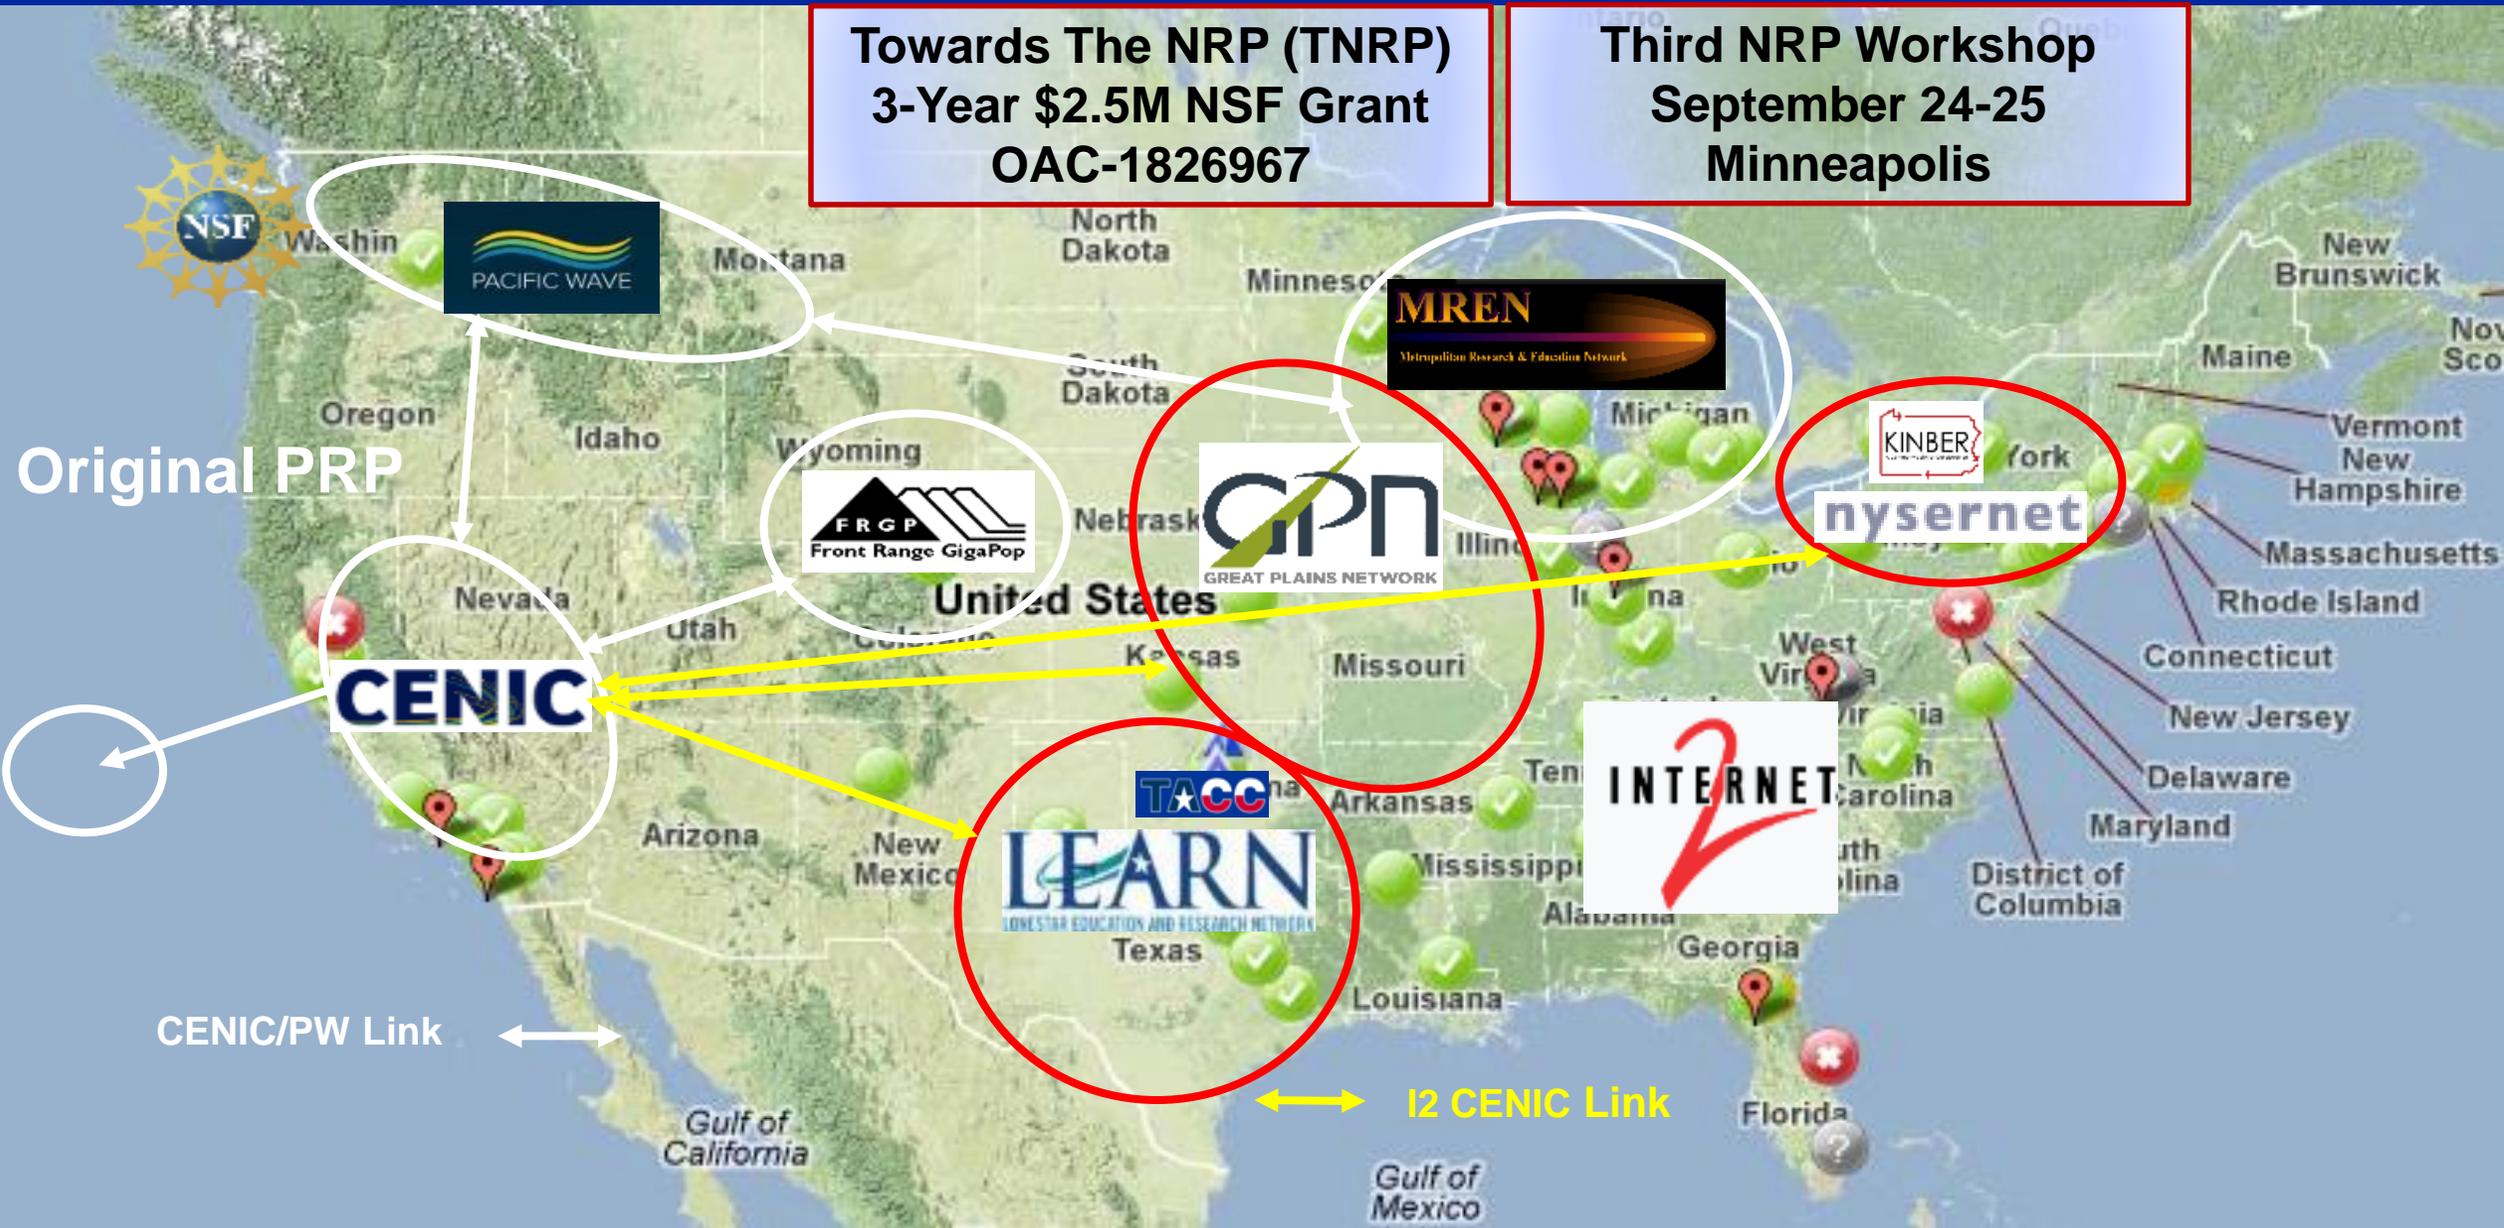
PANELISTS

- Introduction: Jen Schopf, Indiana University / EPOC (Moderator)
- 5 minute overviews
 - Towards the Next Research Platform: Larry Smarr, University of California San Diego / Calit2 / TNRP
 - Engagement and Performance Operations Center: Jennifer Schopf
 - Great Plains Network: James Deaton
 - KINBER: Jennifer Oxenford
 - LEARN: Akbar Kara
- Question Time!

TNRP = PRP (CENIC, PNWGP, FRGP, HI, and MREN) + OSG + ESnet + Quilt + NRP Pilot (I2, KINBER, Learn, GPN, NYSERnet) + ...

**Towards The NRP (TNRP)
3-Year \$2.5M NSF Grant
OAC-1826967**

**Third NRP Workshop
September 24-25
Minneapolis**



Installing FIONAs Across California in Late 2018 and Early 2019 To Enhance User's CPU and GPU Computing, Data Posting, and Data Transfers



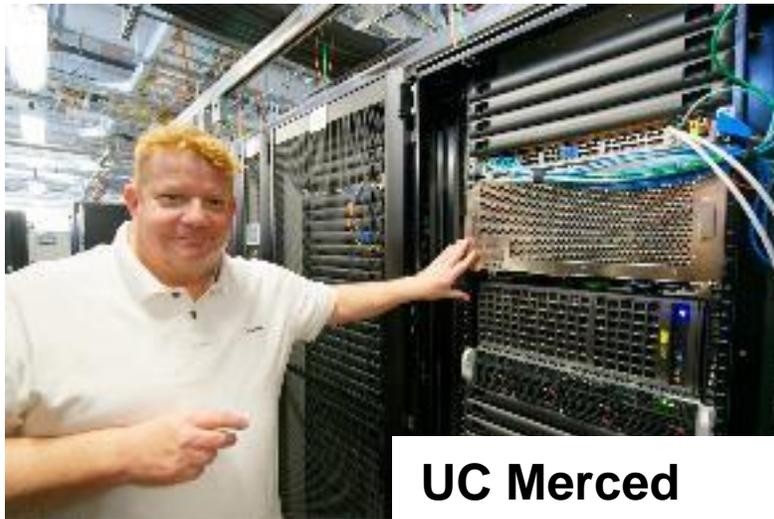
UC Irvine



Stanford



UC Santa Barbara



UC Merced



UC Santa Cruz



UC Riverside

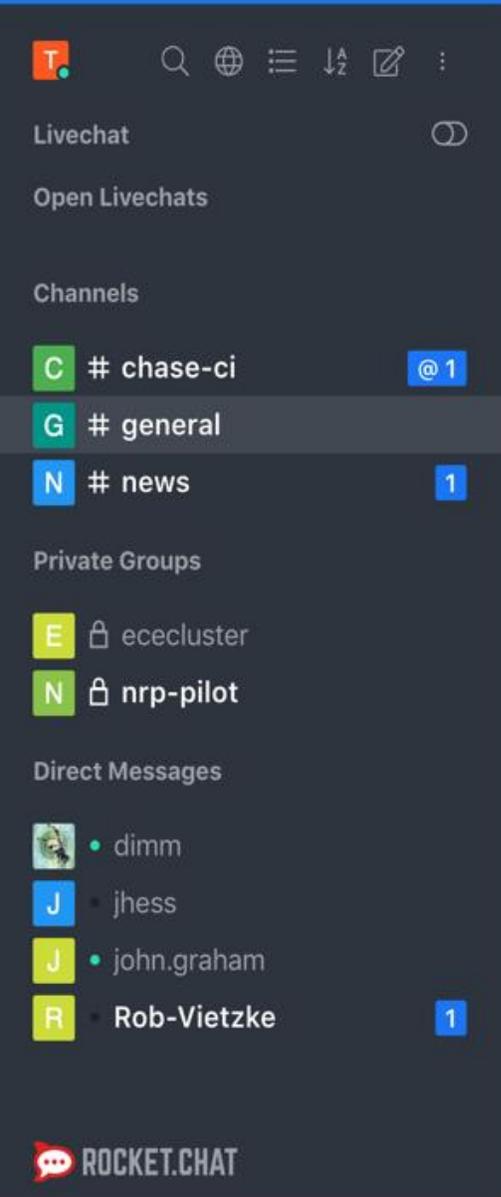
TNRP Patterns of Federation and Sharing

TNRP Nautilus Federation

- **CILogon Federated Authentication**
- **Rook/Ceph/EdgeFS Cloud-Native Storage**
- **Kubernetes Container Orchestration**
- **Network Monitoring perfSONAR MaDDash**
- **Prometheus / Grafana Dashboards**
- **ElastiFlow - EKS Flow Analysis Dashboards**
- **Federated Namespaces**

RocketChat: Helping Users Help Each Other

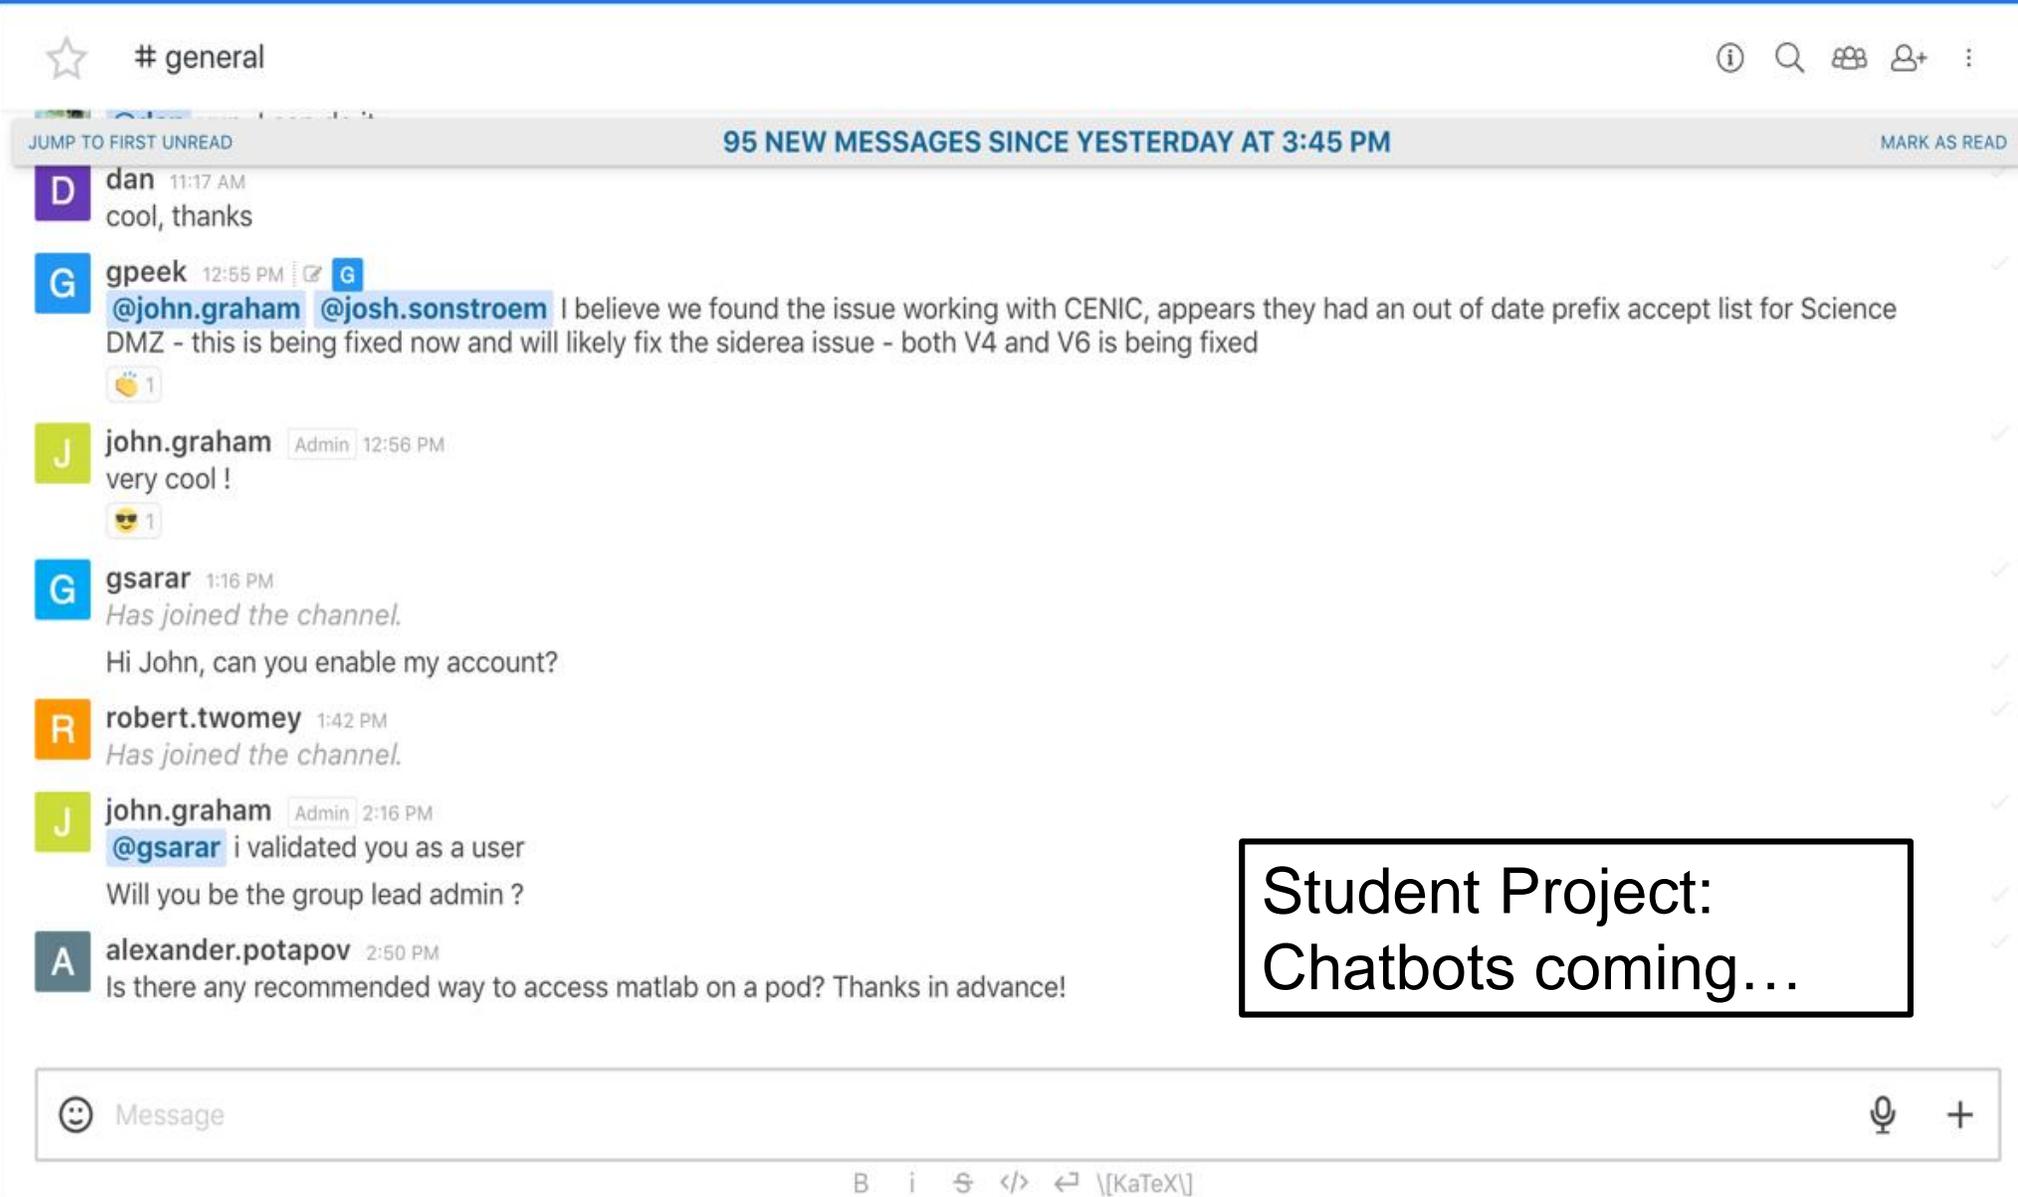
724 Users & 180,000 Messages as of 9/14/19



The sidebar navigation menu is dark-themed and contains the following sections:

- Livechat**: Includes a toggle switch for "Open Livechats".
- Channels**: Lists three channels: "# chase-ci" (1 unread), "# general" (selected), and "# news" (1 unread).
- Private Groups**: Lists two groups: "ececluster" and "nrp-pilot".
- Direct Messages**: Lists four direct messages: "dimm", "jhess", "john.graham", and "Rob-Vietzke" (1 unread).

At the bottom of the sidebar is the "ROCKET.CHAT" logo.



The chat window shows the "# general" channel with a notification for "95 NEW MESSAGES SINCE YESTERDAY AT 3:45 PM". The message history includes:

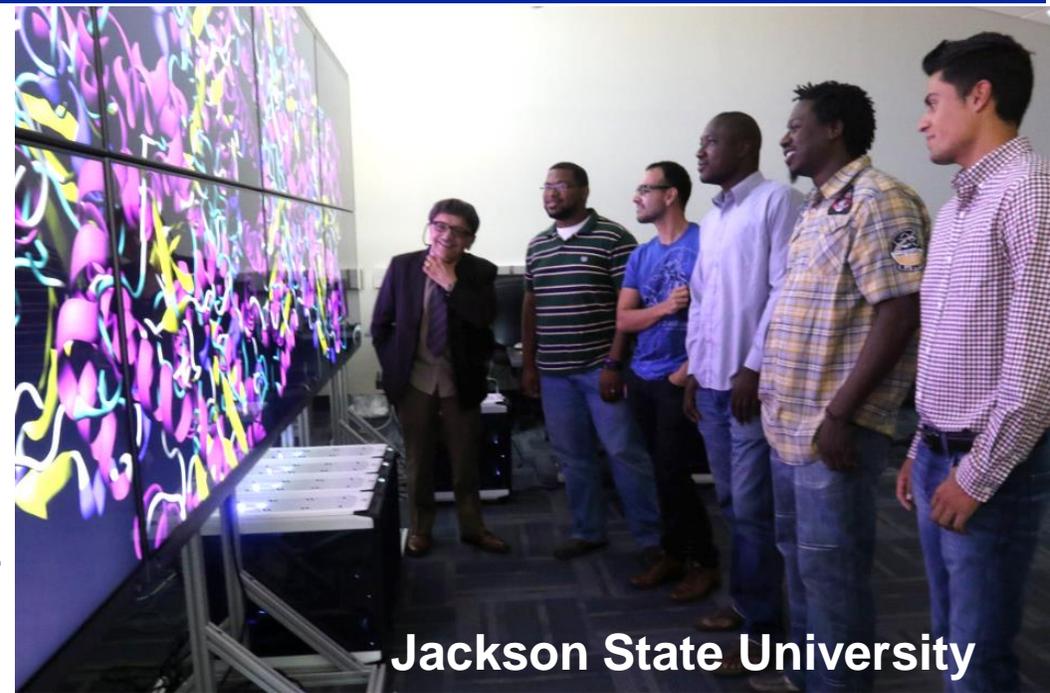
- dan** (11:17 AM): "cool, thanks"
- gpeek** (12:55 PM): "@john.graham @josh.sonstroem I believe we found the issue working with CENIC, appears they had an out of date prefix accept list for Science DMZ - this is being fixed now and will likely fix the siderea issue - both V4 and V6 is being fixed"
- john.graham** (Admin, 12:56 PM): "very cool!"
- gsarar** (1:16 PM): "Has joined the channel. Hi John, can you enable my account?"
- robert.twomey** (1:42 PM): "Has joined the channel."
- john.graham** (Admin, 2:16 PM): "@gsarar i validated you as a user Will you be the group lead admin?"
- alexander.potapov** (2:50 PM): "Is there any recommended way to access matlab on a pod? Thanks in advance!"

The bottom of the chat window features a message input field with a "Message" placeholder, a microphone icon, and a plus sign for additional actions. A rich text editor toolbar is visible at the very bottom.

Student Project:
Chatbots coming...

PRP/TNRP Actively Develops Diversity

- **Grants**
 - **3 Female co-PIs**
 - **1 Hispanic co-PI**
- **Campuses**
 - **8 Minority-Serving Institutions in PRP/CHASE-CI**
- **Workshops**
 - **NRPII Workshop Steering Committee 80% Female**
 - **Multiple MSI, EPSCoR Focused Workshops**



Presenting
FIONettes



EPOC

Engagement and Performance
Operations Center

The Engagement and Performance Operations Center (EPOC)

Dr. Jennifer M. Schopf
PI, EPOC

Indiana University International Networks

Jason Zurawski
Co-PI, EPOC

ESnet / Lawrence Berkeley National
Laboratory



ESnet

ENERGY SCIENCES NETWORK



INDIANA UNIVERSITY

Engagement and Performance Operations Center

- Joint project between Indiana University and ESnet
 - co-PI Jent (IU GlobalNOC) and Zurawski (ESnet)
- Part of CC* program for domestic science support
 - Program Officer: Kevin Thompson
- Award #1826994, \$3.5M over 3 years
- Partnerships with regional, infrastructure, and science communities that span the NSF and DOE continuum of funding

Understanding End-to-End Performance is Hard

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

Roadside Assistance Process

- “This file transfer worked last week, but it doesn’t anymore?”
 - Think of this like a flat tire, crash repair
 - Anyone can submit
- Contact epoc@iu.edu
 - Within 24 hours, gets triaged
 - Some initial investigation to verify the issues
 - A Case Manager and Lead Engineer are assigned
 - Shareable infrastructure set up
- Centralization of Researcher Interactions

EPOC Deep Dives

- Think of this as regular maintenance, oil change, or planning to buy a new car
- Based on seminal work by ESnet to develop Scientific Case Studies
 - Walk through science workflow with the actual scientists
 - Way to understand needs and planning
- Often identifies issues that have **nothing** to do with networks, and everything to do with sociology

Network Instrumentation using NetSage

- Performance and measurement are 2 sides of a coin
- Common basic measurement data is the first step to understanding performance issues
 - E.g. Global perfSONAR Deployment, <http://my.es.net>
- NetSage framework
 - SNMP, perfSONAR, Flow, Tstat Data
 - Grafana-based dashboards to visualize performance
- <http://gpn.netsage.global>
- <http://ilight.netsage.global>



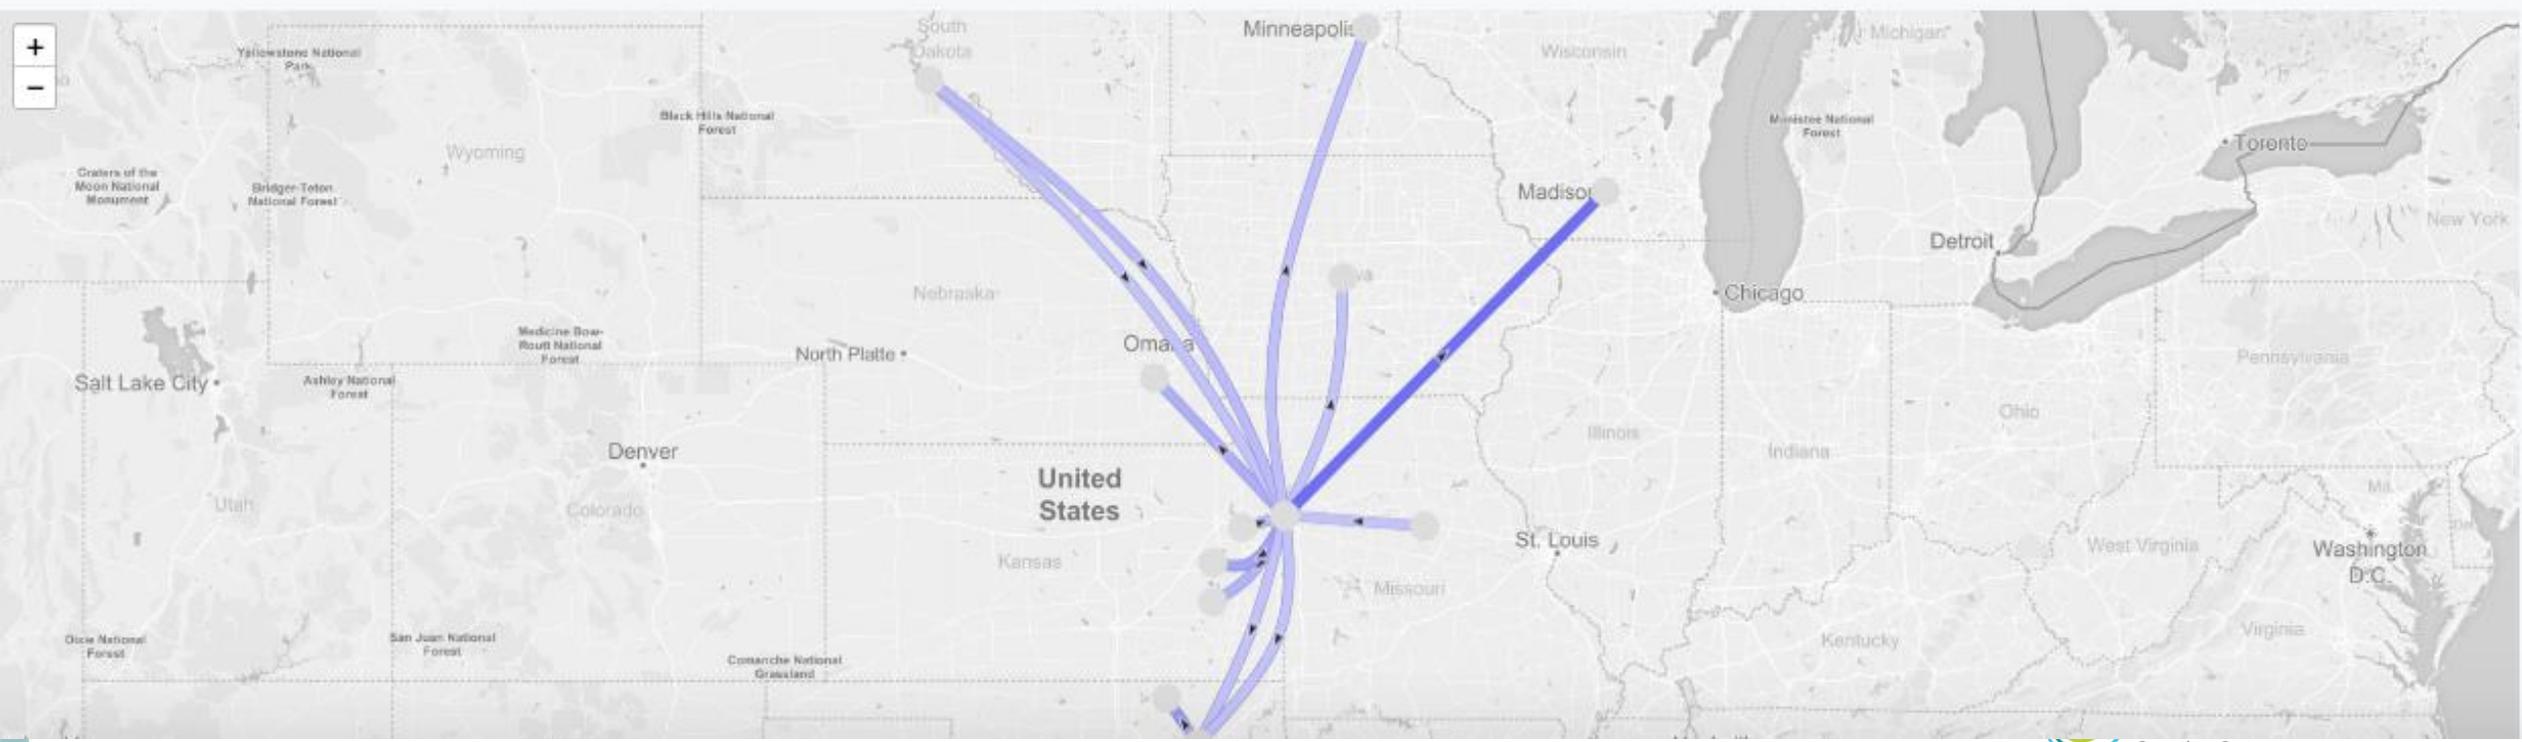
Bandwidth Dashboard



Single Link Max
39.77 Gbps

Average Across All Links
3.70 Gbps

Total Transferred
79.9 TB



Bandwidth patterns

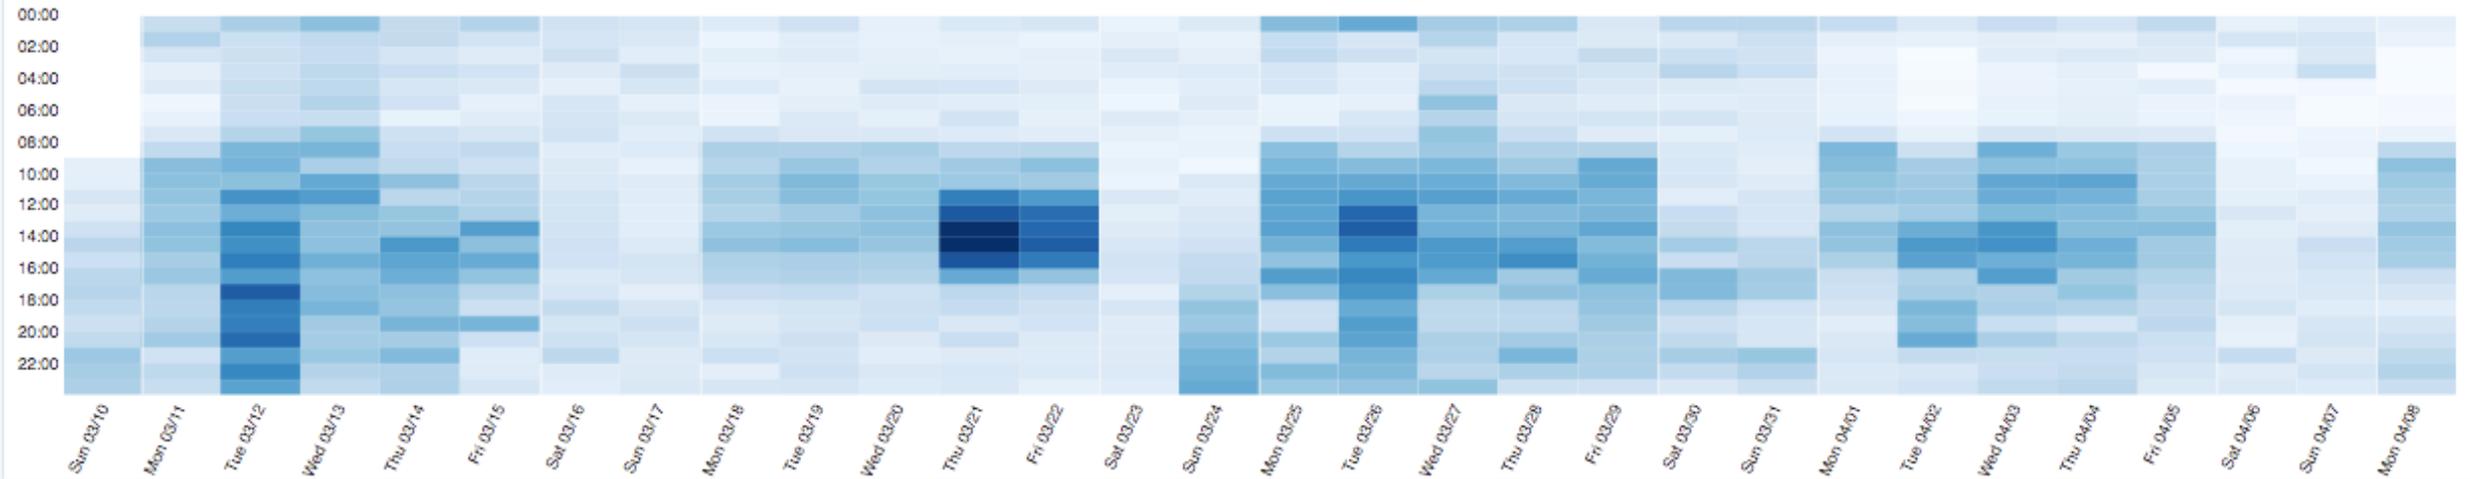


Bandwidth Patterns ▾

🔗 🖨️ 🕒 Last 30 days 🔍 ↻

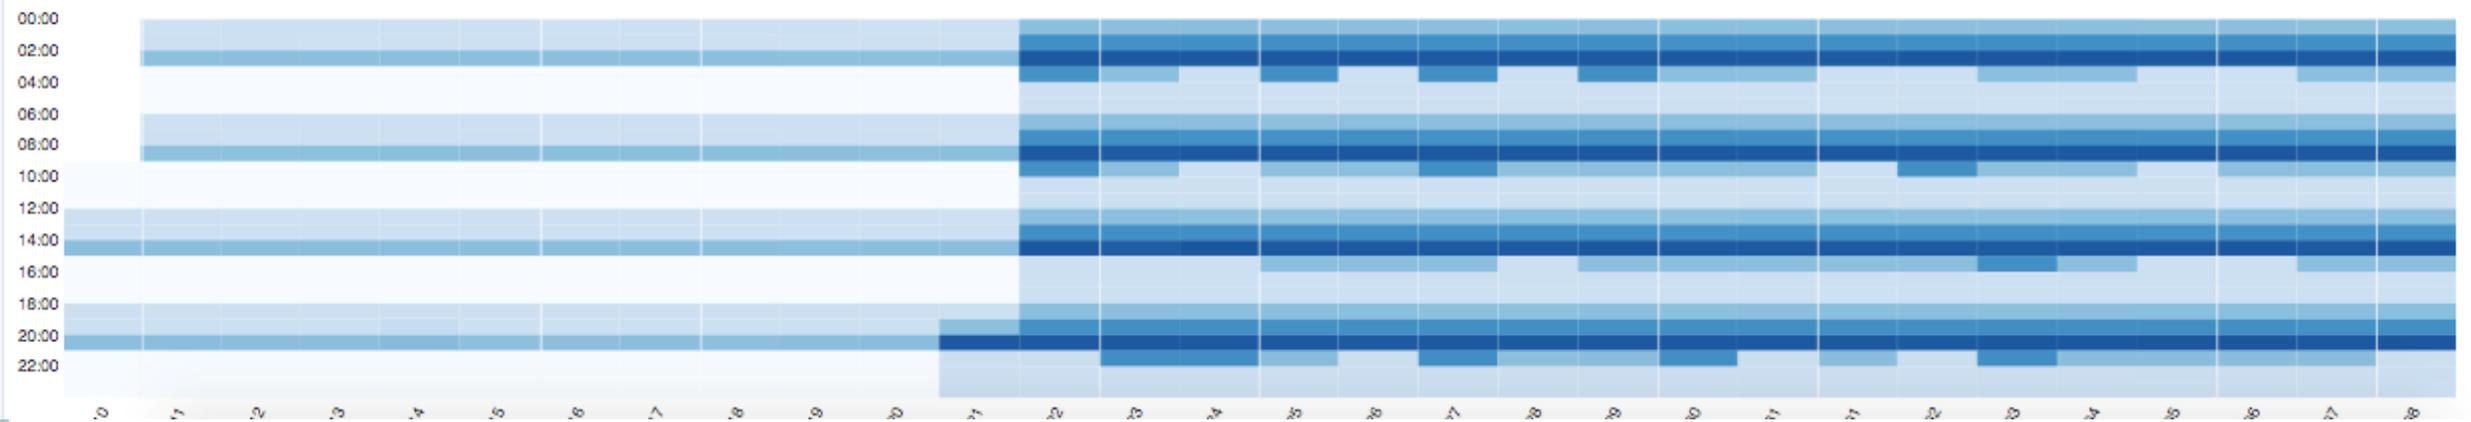
GPN: University of Minnesota 10GE (input)

🕒 Last 30 days

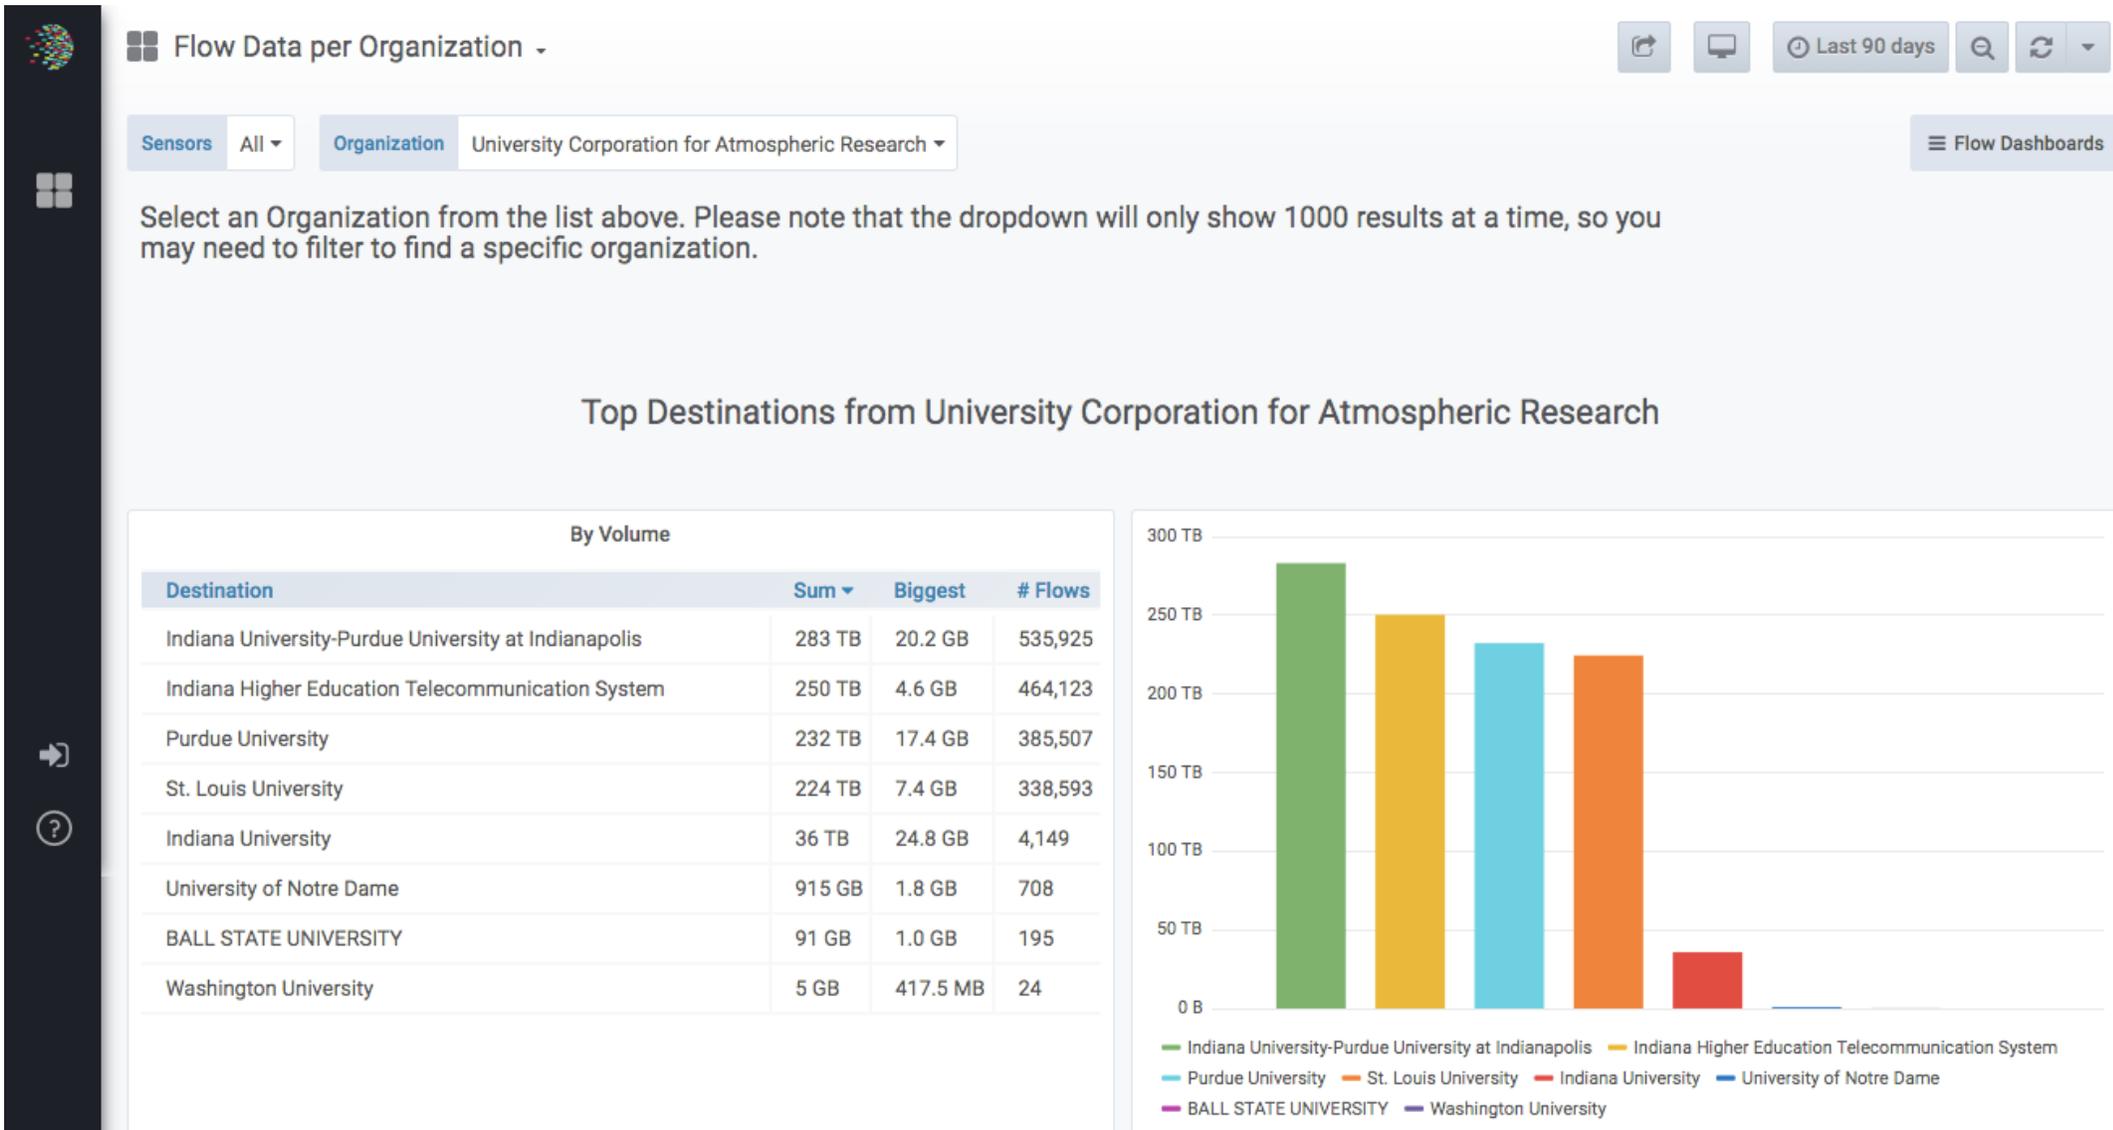


GPN: University of Missouri 100GE (input) ▾

🕒 Last 30 days



Flow data by Organization- iLight and NCAR



What is a “Service-in-a-Box”?

- Basic idea:
 - Only large facilities with dedicated funding can afford the time/effort to design/install/operate/maintain a dedicated science infrastructure
 - Ameliorate the costs of design/install at a higher level (e.g. regional network).
 - Create infrastructure that can be delivered as a service
 - Operation can be local or regional (offer flexibility based on the environment and resources available)
 - ***Develop a business model that facilitates cost recovery and upgrade schedules***

Training

- Follow on to OIN (<http://oinworkshop.com>) series that reached over 750 people in the NSF/DOE funding space during the 3 year operational period
- Hands on perfSONAR sessions
 - Especially for small nodes, includes file transfer tests
- How to do an Application Deep Dive
 - Also known as “How to talk to Scientists”
- DMZ/DTN Set Up
- To request – send mail to epoc@iu.edu
 - include “Training Request: in the subject line

Take Aways

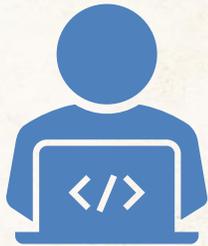
- EPOC is an NSF-funded operations center to help scale science engagement and problem resolution
- Single point of contact to help with end-to-end performance issues
 - epoc@iu.edu
- More about EPOC:
 - <http://epoc.global>
- Jennifer Schopf, jmschopf@iu.edu
- Jason Zurawski, zurawski@es.net

EPOC in Texas

Akbar Kara, Chief Technology Officer

September 25, 2019

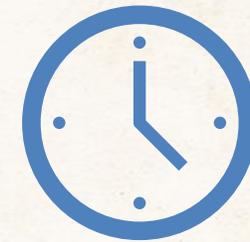
Overview – Trinity University



6 IT Staff



5 Researchers

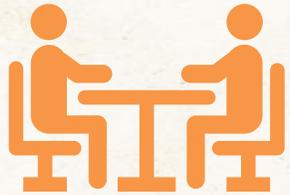


8 hour engagement



**Classical Studies, Chemistry, Geosciences,
Computer Science, Physics & Astronomy**

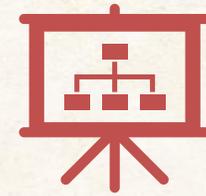
Impact



**Bringing IT & Faculty
together**



Ownership



**Assessment of existing
infrastructure**



Interpretation of needs



Recommended Solutions

Interest from other Texas Universities



Texas Tech University

Lubbock - West Texas

Baylor University

Waco - Central Texas





How Regional Partnerships with National Performance Engineering and Outreach Initiatives are Enabling Science

Jennifer Oxenford, Director, Community Engagement, KINBER
Co-located NSF CC* and CICI PI Workshop/TNRP/Quilt Fall
Member Meeting
Minneapolis, MN
Wed. Sept. 25, 2019

KINBER Background/Overview

- **Keystone Initiative for Network Based Education and Research (KINBER)** is Pennsylvania's statewide research, education, and community network.
- **Non-profit, member based** organization that provides a variety of **next generation broadband connectivity** solutions and services to community anchor organizations in PA.
- Formed in 2009 to receive capital funding through a **\$99.6 million NTIA grant, with over \$29M matching funds.**
- Owns and operates **1800+ mile PennREN** fiber optic infrastructure in Pennsylvania.
- One of **43 Research and Education Networks nationwide.**
- KINBER offers **network connectivity** to its advanced **PennREN network** and **trusted technology solutions** to its member community.

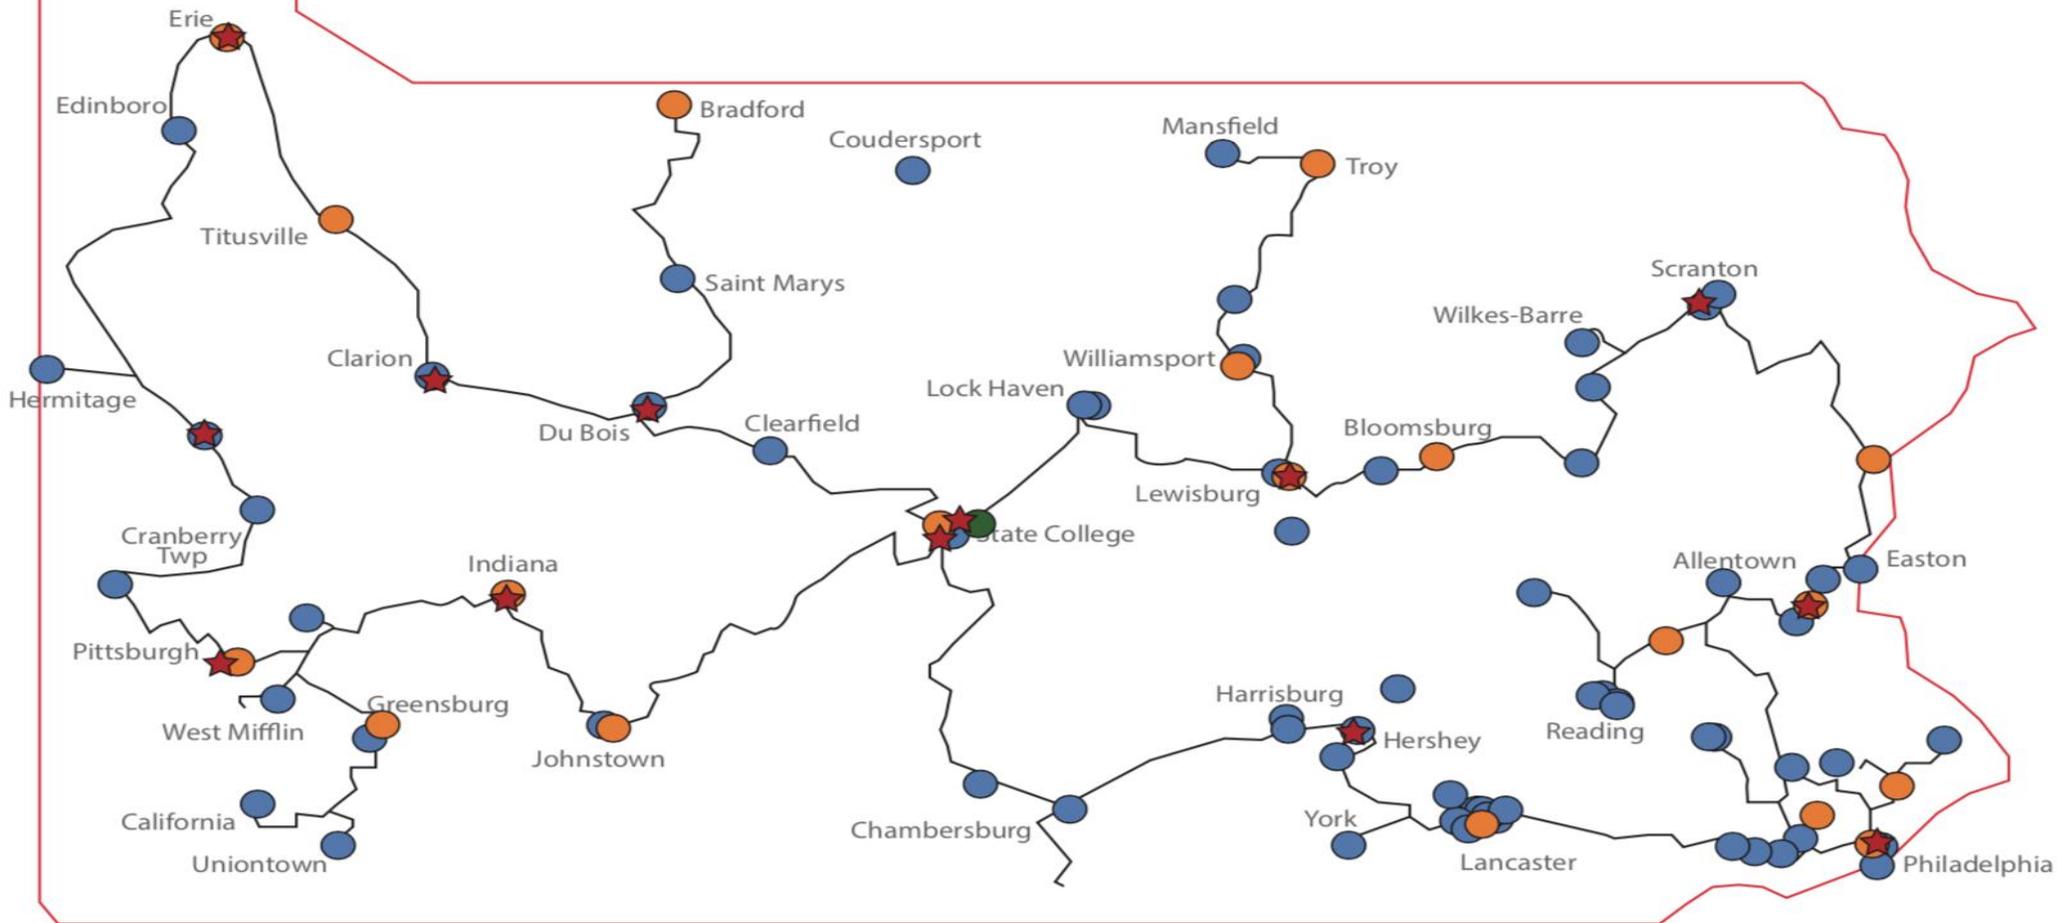


CELEBRATING 10 YEARS
*of Pennsylvania's Statewide Research,
Education, and Community Network*

www.kinber.org

KINBER Network Map

KINBER, the Keystone Initiative for Network Based Education and Research, is a trusted technology partner that provides a strategic and competitive advantage to Pennsylvania-based organizations through high-speed broadband connectivity, collaboration, and innovative use of digital technologies.



- 1 GE Connection
- 10 GE Connection
- 100 GE Connection
- ★ Service Nodes



KINBER Support for Enabling Science

Since 2014 KINBER has provided education and training on enabling science and supporting research through webinars, workshops, virtual trainings, pilots, and more.

- OIN workshop in 2016 at Millersville University
- Preparing and Submitting your NSF CC* Grant Proposal webinar series 2016-2019
- PerfSONAR and Network Performance workshops at KINBER annual conferences
- Webinars and trainings on Globus, ESnet Science Engagement, XSEDE Campus Champions, OSG, etc.
- Virtual Office Hours for Colleges/Universities
- Support through KINBER Research Engagement Advisory Committee

KINBER Support for Enabling Science

NSF Campus Cyberinfrastructure CC* grants in PA

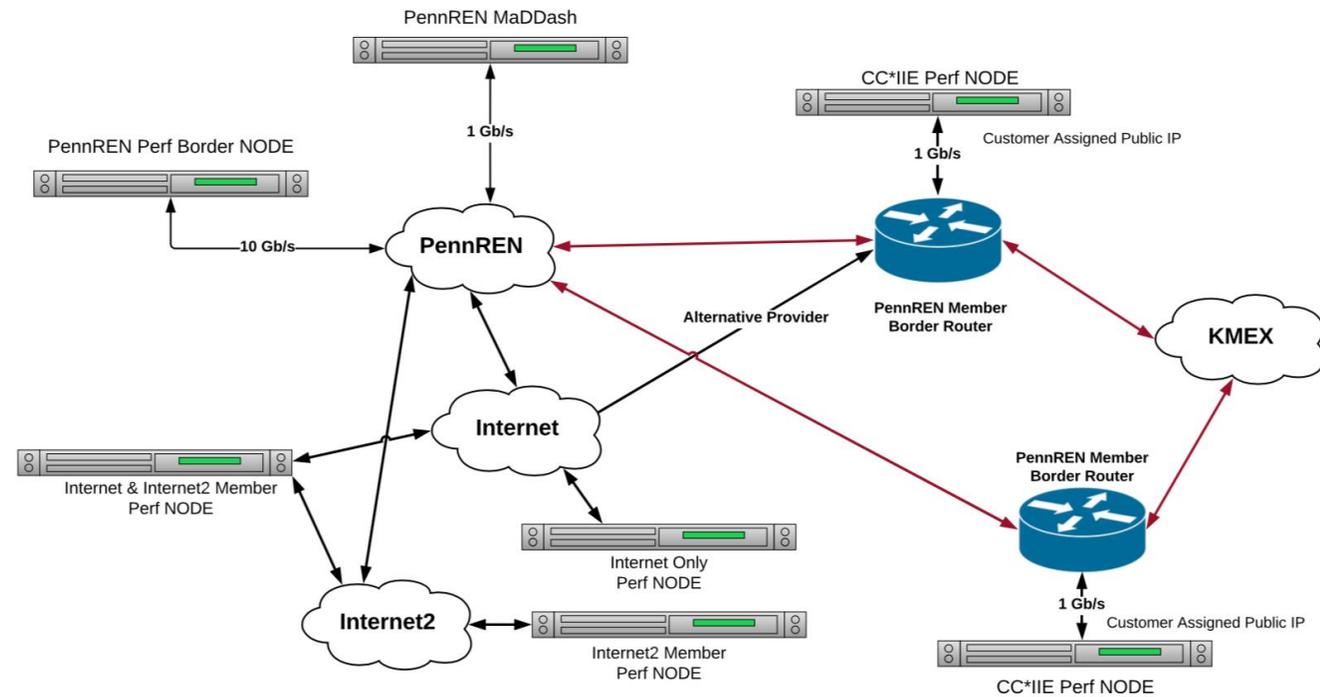
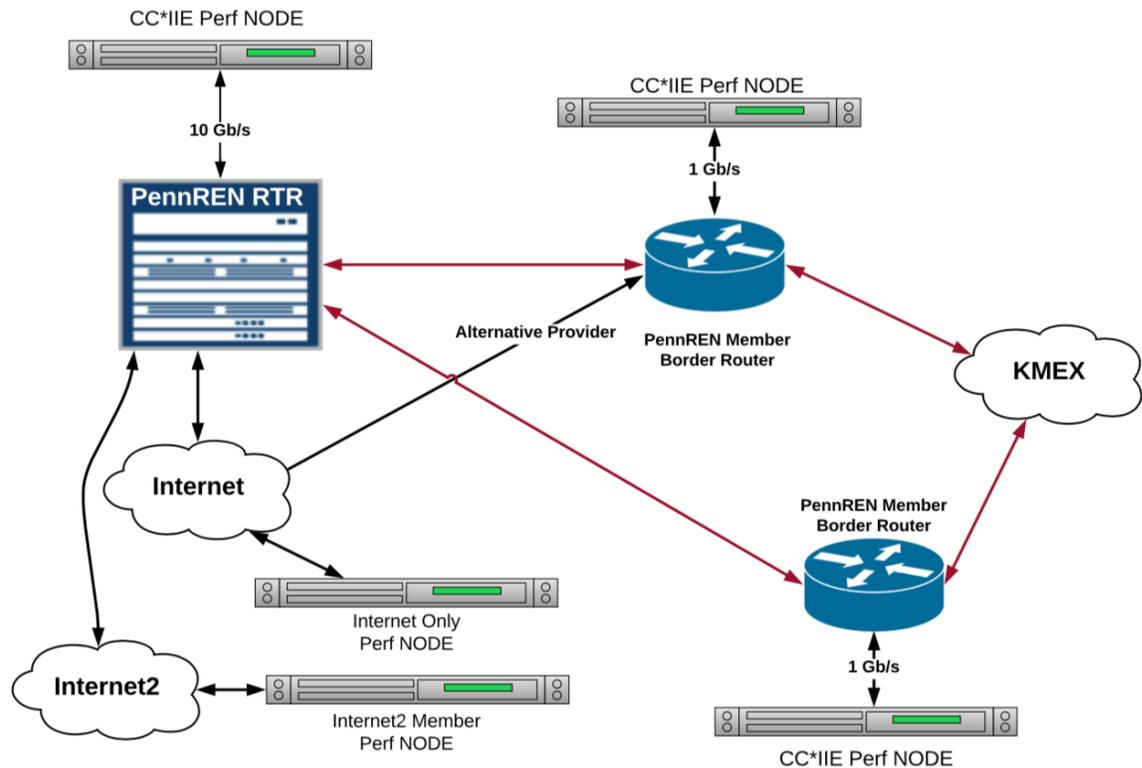
- Franklin & Marshall College 2015 CC*DNI Campus Design: Building a State of the Art Research Network at Franklin and Marshall College; \$350K; 2019 CC* Compute: Building a state-of-the-art campus compute resource at Franklin & Marshall College \$400K
- Bucknell University 2017 CC* Network Design: The Bucknell Science DMZ Network Design and Implementation \$399,200K
- Arcadia University 2018 CC* Transforming Arcadia's Networking Capability, Enhancing for Innovation to Grow Research Leaders in a Technology-driven World \$352K
- Duquesne University 2019 CC* Networking: A High-Performance Science DMZ and Dedicated Research Network for Duquesne University \$462K

KINBER CC* perfSONAR Activities

- KINBER NSF CC*IIE Region Accelerating the Adoption of Campus Cyberinfrastructure Technologies in Pennsylvania Grant (Award 1440699)
- As part of this grant, KINBER offered a perfSONAR pilot program from 2015 – 2018 to several schools in PA.
- Pilot provided pre-configured perfSONAR platforms to six campuses in PA in order to understand if pre-configured platforms lead to faster integration into the campus infrastructure.
- Participating campuses included: Franklin & Marshall College, Swarthmore College, CMU, Villanova, Harrisburg Area Community College, and The Hill School
- Representative campuses included K12, R1, and small and under-resourced schools.
- KINBER perfSONAR Pilot provided initial key findings on perfSONAR deployment in region.

KINBER perfSONAR Activities

KINBER PennREN perfSONAR Topology



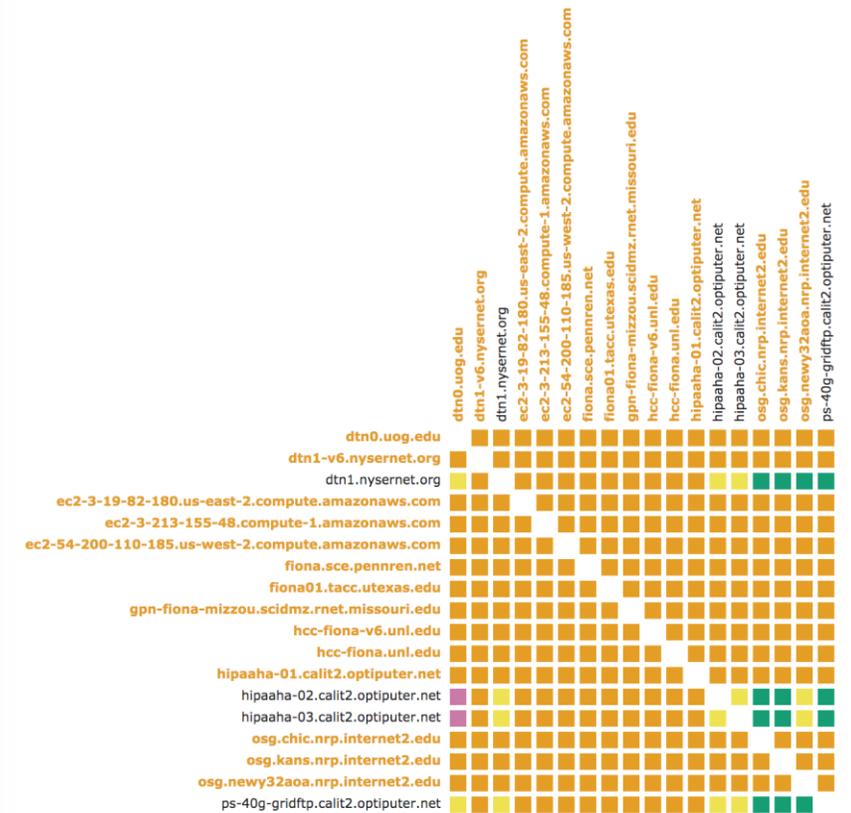
Enhanced Support with PRP/NRP

KINBER is part of the PRP/NRP pilot announced in 2018 at Internet2 Global Summit.

KINBER staff actively participate in NRP activities including PRP/NRP Pilot Engagement team, PRP Engineering group, and Scaling the National Research Platform Communications team.

KINBER has deployed PRP infrastructure and tools including a FIONA node.

KINBER continues to seek and share opportunities for regional awareness and training on PRP/NRP infrastructures and activities to enable and support science in PA.



NRP Mesh - Throughput NRP_GridFTP - Throughput

KINBER and Eastern Regional Network (ERN)

Understand and support end-to-end performance and data movement in support of research in local, regional and national environments.

Project - NSF Funded DIBBS VDC

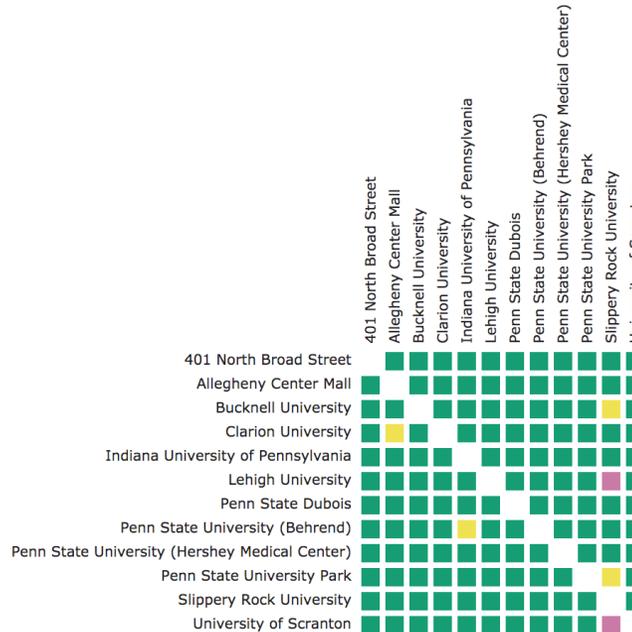
NSF DIBBs Virtual Data Collaboratory - Throughput, full-mesh - Throughput

■ Throughput >= 1Gbps
 ■ Throughput < 1Gbps
 ■ Throughput <= .5Gbps
 ■ Unable to find test data
 ■ Check has not run yet



Pennsylvania - PennREN

PennREN Core Mesh - PennREN CORE 1G Throughput (KMEX)

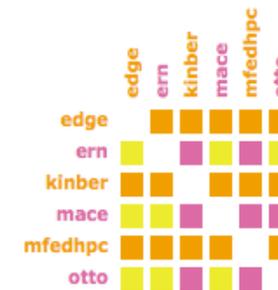


ERN

Dashboard 1 - Dashboard 1 - throughput_task - Throu

■ Throughput >= 1Gbps
 ■ Throughput < 1Gbps
 ■ Throughput <= .5Gbps

⚠ Found a total of 7 problems involving 6 hosts in the grid



Enhanced Support with EPOC Team – Outreach/Training

EPOC staff Jason Zurawski and Hans Addleman participated in KINBERCON April 1-3, 2019 in Philadelphia. Gave general session talk *EPOC (Engagement Performance and Operations Center) Services and Activities Supporting Research in PA*.

EPOC staff Jason Zurawski and Hans Addleman offered a half-day workshop *Determining Technology Requirements for Scientific Innovation*.

In addition to the talk and workshop they led a Deep Dive with Arcadia University as part of the workshop.

Result of the workshop and deep dive generated a White Paper on Arcadia University Bioinformatics Deep Dive <https://escholarship.org/uc/item/1196z33x>

Enhanced Support with EPOC Team – Arcadia Deep Dive

The Arcadia University bioinformatics project was one of the science drivers for the successful NSF Campus Cyberinfrastructure award entitled “Transforming Arcadia’s Networking Capability, Enhancing for Innovation to Grow Research Leaders in a Technology-driven World.”

3-year CC* award for \$352,500 started in 2018. The project is creating a Science DMZ with a data transmission network capable of 10Gbps connectivity (more than 10 times faster than current speeds) to KINBER and national research networks.

Side meetings were held among EPOC staff, KINBER staff, Arcadia, and Penn State University staff to determine future collaboration opportunities and engagement.

Enhanced Support with EPOC Team - NetSage

- KINBER is working with closely with EPOC team on setting up our NetSage monitoring suite deployment in PA for proactive discovery and resolution of performance issues via network analysis services and application analysis.
- Expected to have this set up by end of Oct. 2019. Working with GlobalNOC to obtain SNMP data.
- Goal moving forward to set up Net Sage using flow data as well.
- Franklin and Marshall College interest in setting up NetSage.
- KINBER committed to deploying and using NetSage. Also plan/hope to work with EPOC, GPN and others towards possible containerized version of the extra flow data.

Enhanced Support with EPOC Team - perfSONAR and DMZ

- KINBER in middle of perfSONAR Managed Service project with EPOC. Looking to provide “perfSONAR in a Box” to select campuses.
- KINBER member Arcadia University is interested piloting the “perfSONAR in a box” deployment with EPOC as initial KINBER perfSONAR pilot site, to launch, test, and refine the model.
- Goal to scale the “perfSONAR in a Box” model to additional colleges and universities in PA and then grow to add additional offerings in partnership with EPOC.
- Additional perfSONAR, DMZ, and other managed service interest and opportunities being explored among KINBER community.

Thank you

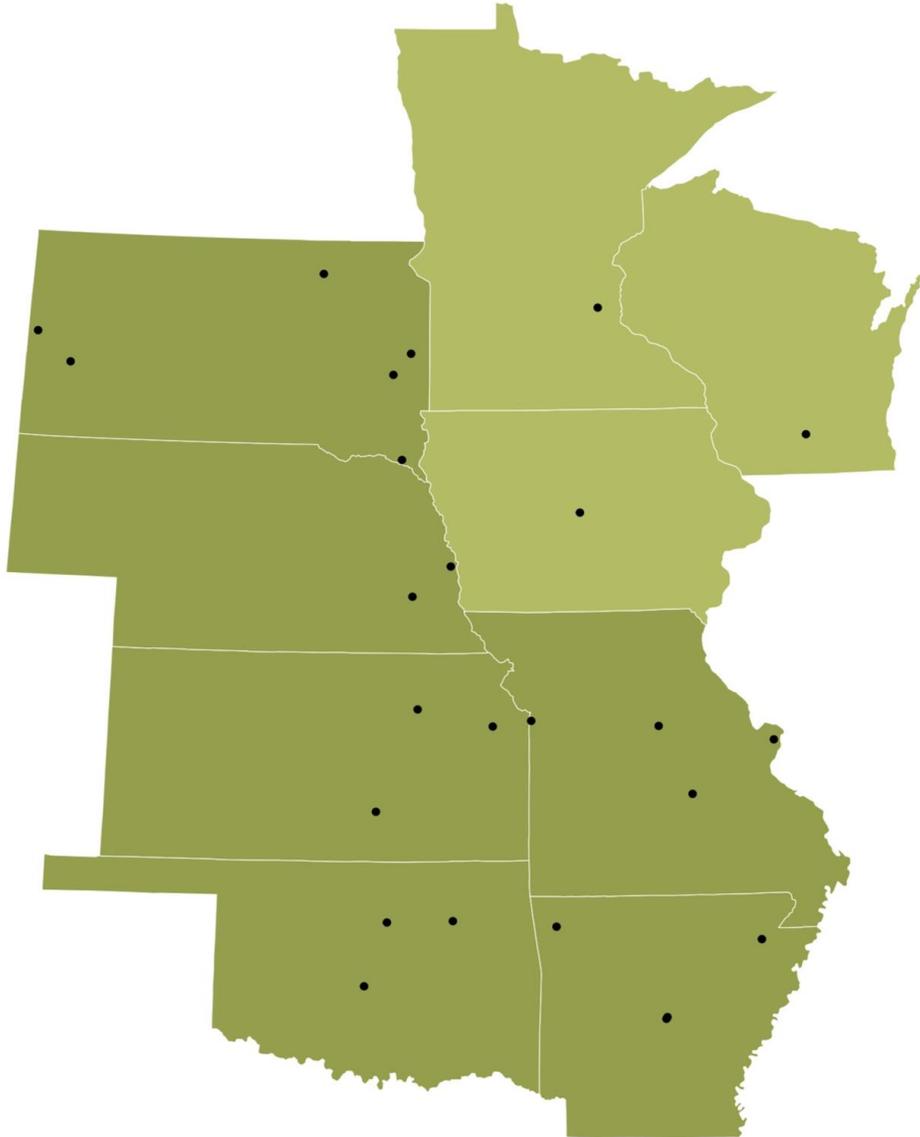
Jennifer Oxenford
Director Community Engagement
joxenford@kinber.org



EPOC & TNRP impact within GPN

NSF / NRP / The Quilt
Minneapolis, MN
2019

James Deaton
Executive Director
jed@greatplains.net



- The Great Plains Network (GPN) is a non-profit consortium aggregating networks through GigaPoP connections while advocating research on behalf of universities and community innovators across the Midwest and Great Plains who seek collaboration, cyberinfrastructure and support for big data and big ideas, at the speed of the modern Internet.
- Over two dozen universities
 - Across 9 states
 - More than 20 years of collaborating in research and education networking



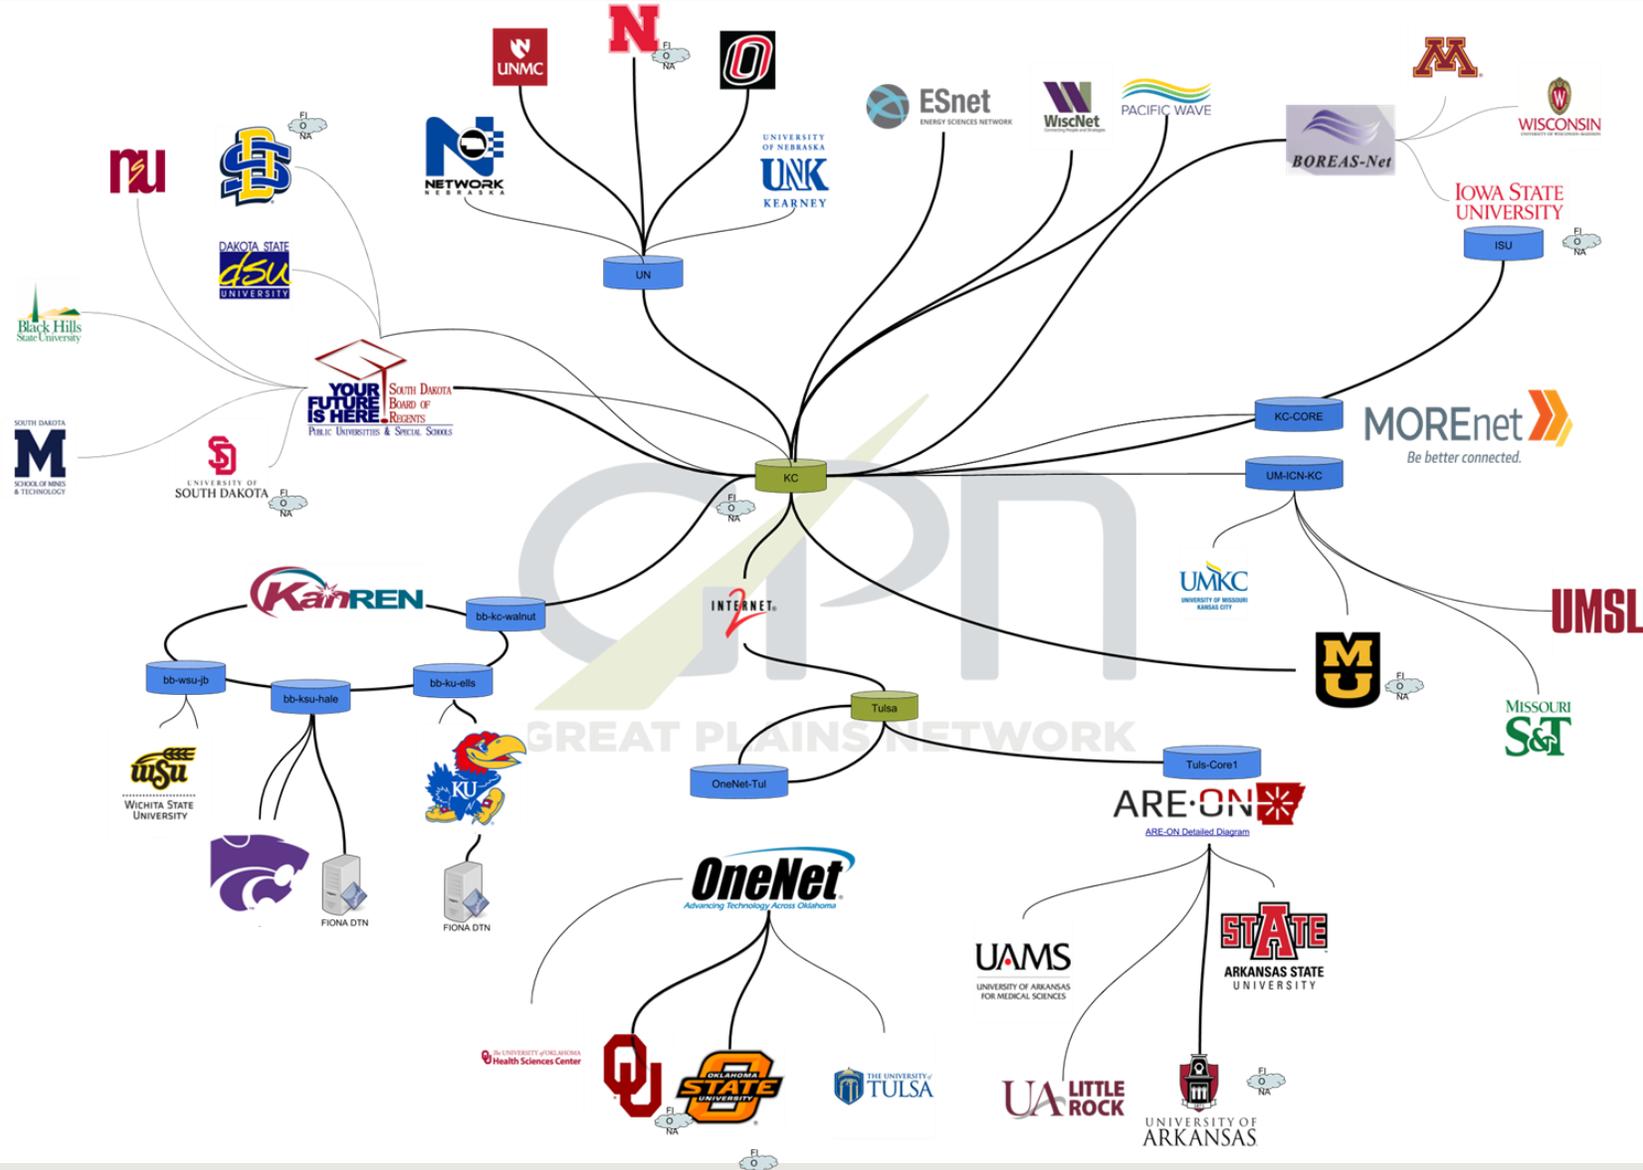
Connected via 6 state networks and a RON

- ARE-ON
- KanREN
- MOREnet
- Network Nebraska
- OneNet
- SD-REED

- BOREAS

Goals with NetSage and More Broadly via work with EPOC

- Deeper understanding of R&E data movement across the region
- Engaging campus IT and state networks in both passive and active approaches of improving data movement
- Increasing awareness of the rich set of science drivers within the region



EPOC Engagements and Deployments

- Workshop: EPOC Science Engagement: Bridging the Technology Gap between Information Technology and the Research Community
- What are GPN Folks Doing Internationally
- BoF: Ask the CI Engineer
- NetSage

Single Link Max
69.76 Gbps

Average Across All Links
4.19 Gbps

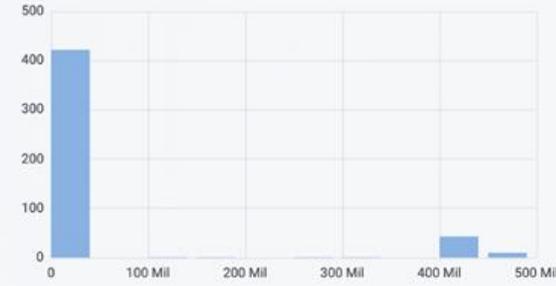
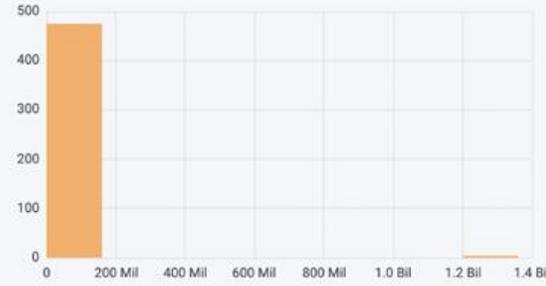
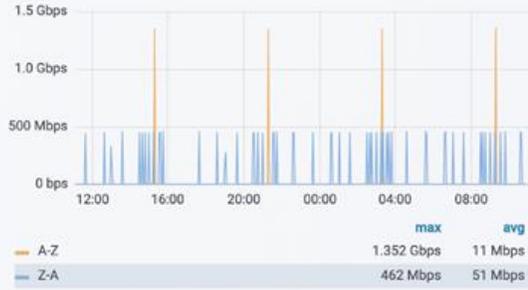
Total Transferred
5.42 PB





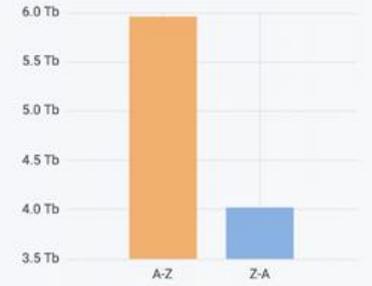
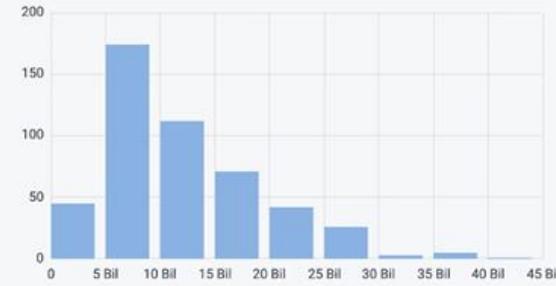
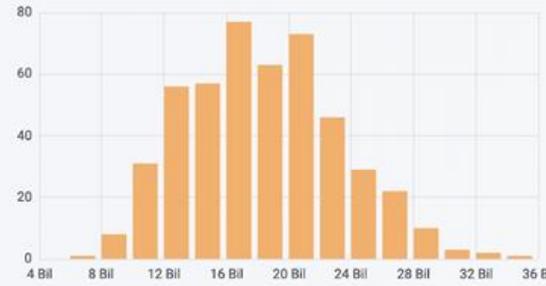
Bandwidth Dashboard -

GPN: University of Missouri 100GE



Last 24 hours Refresh every 1d

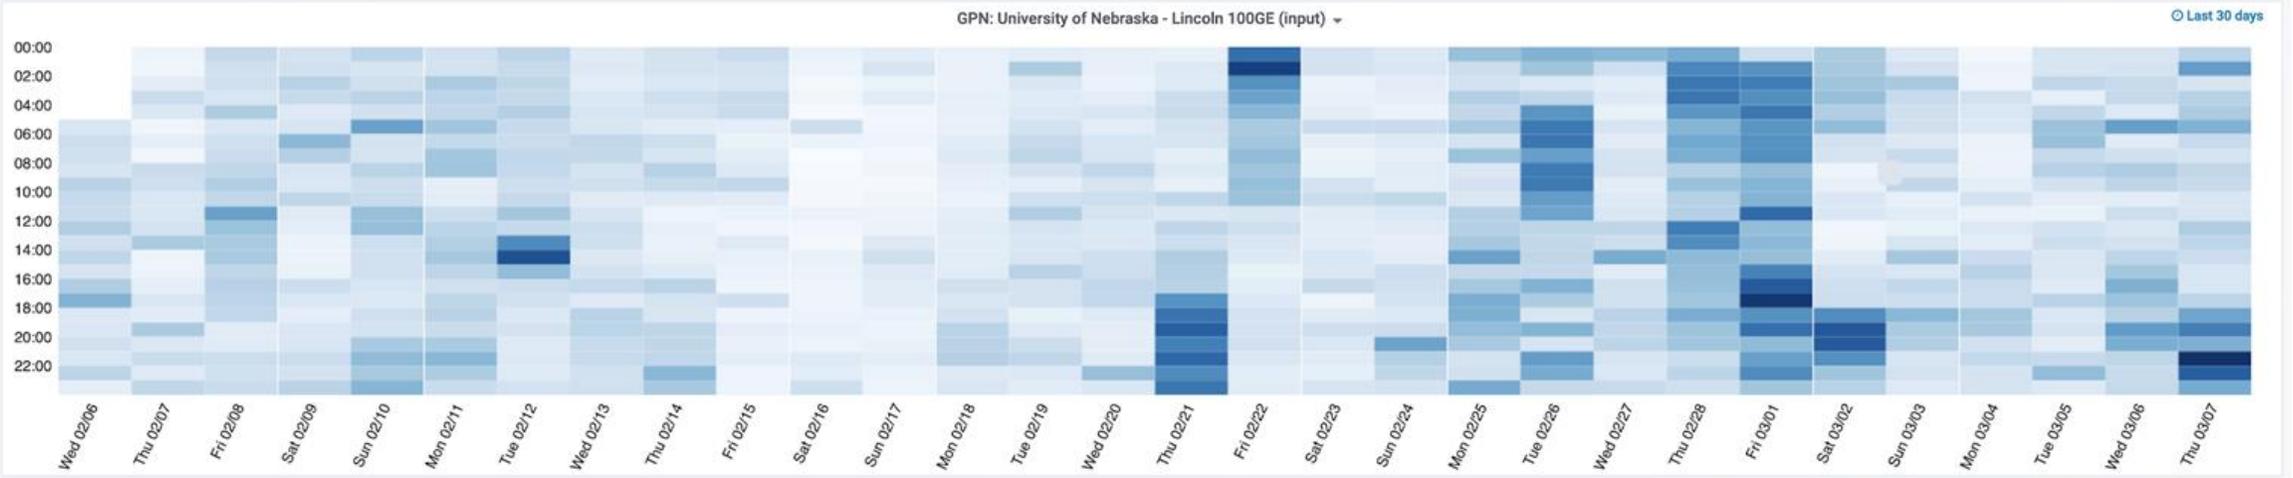
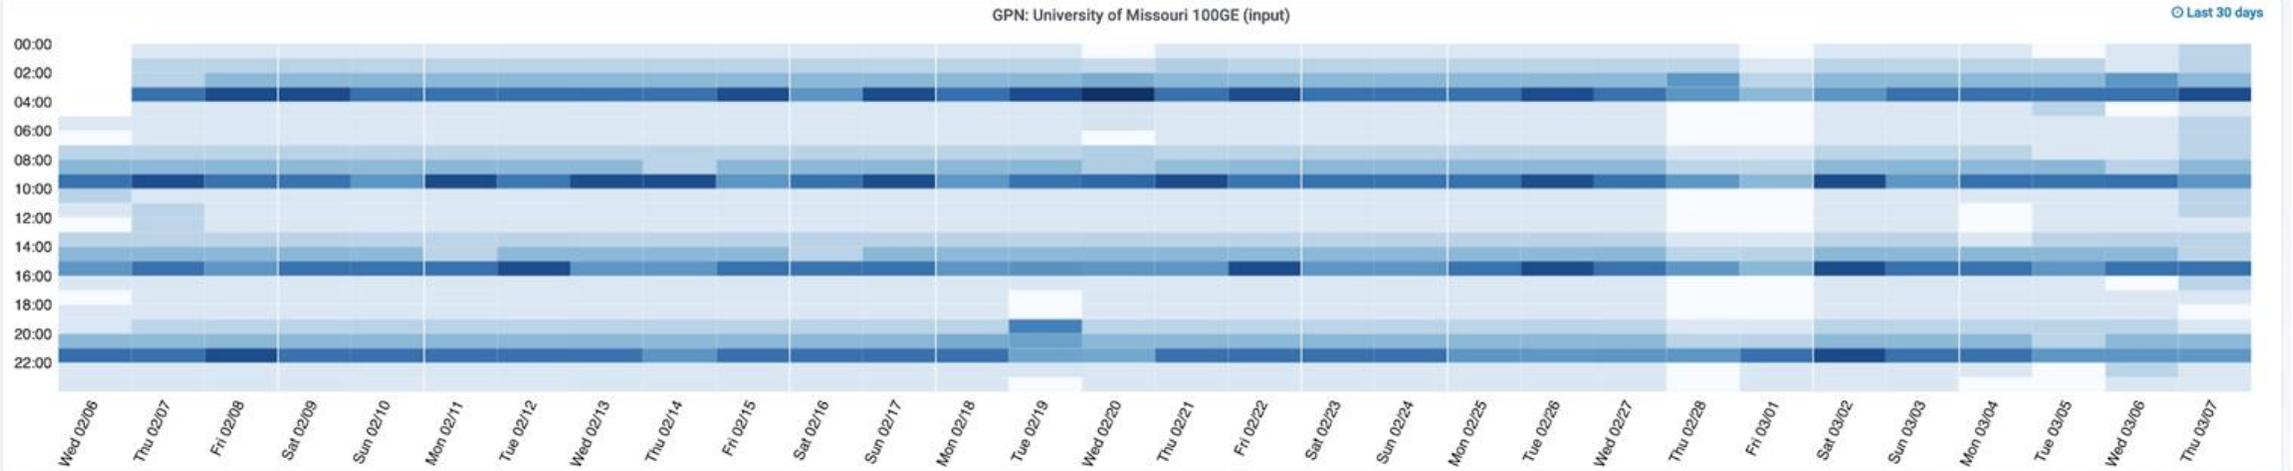
GPN: University of Nebraska - Lincoln 100GE



GPN: WiscNet 10GE

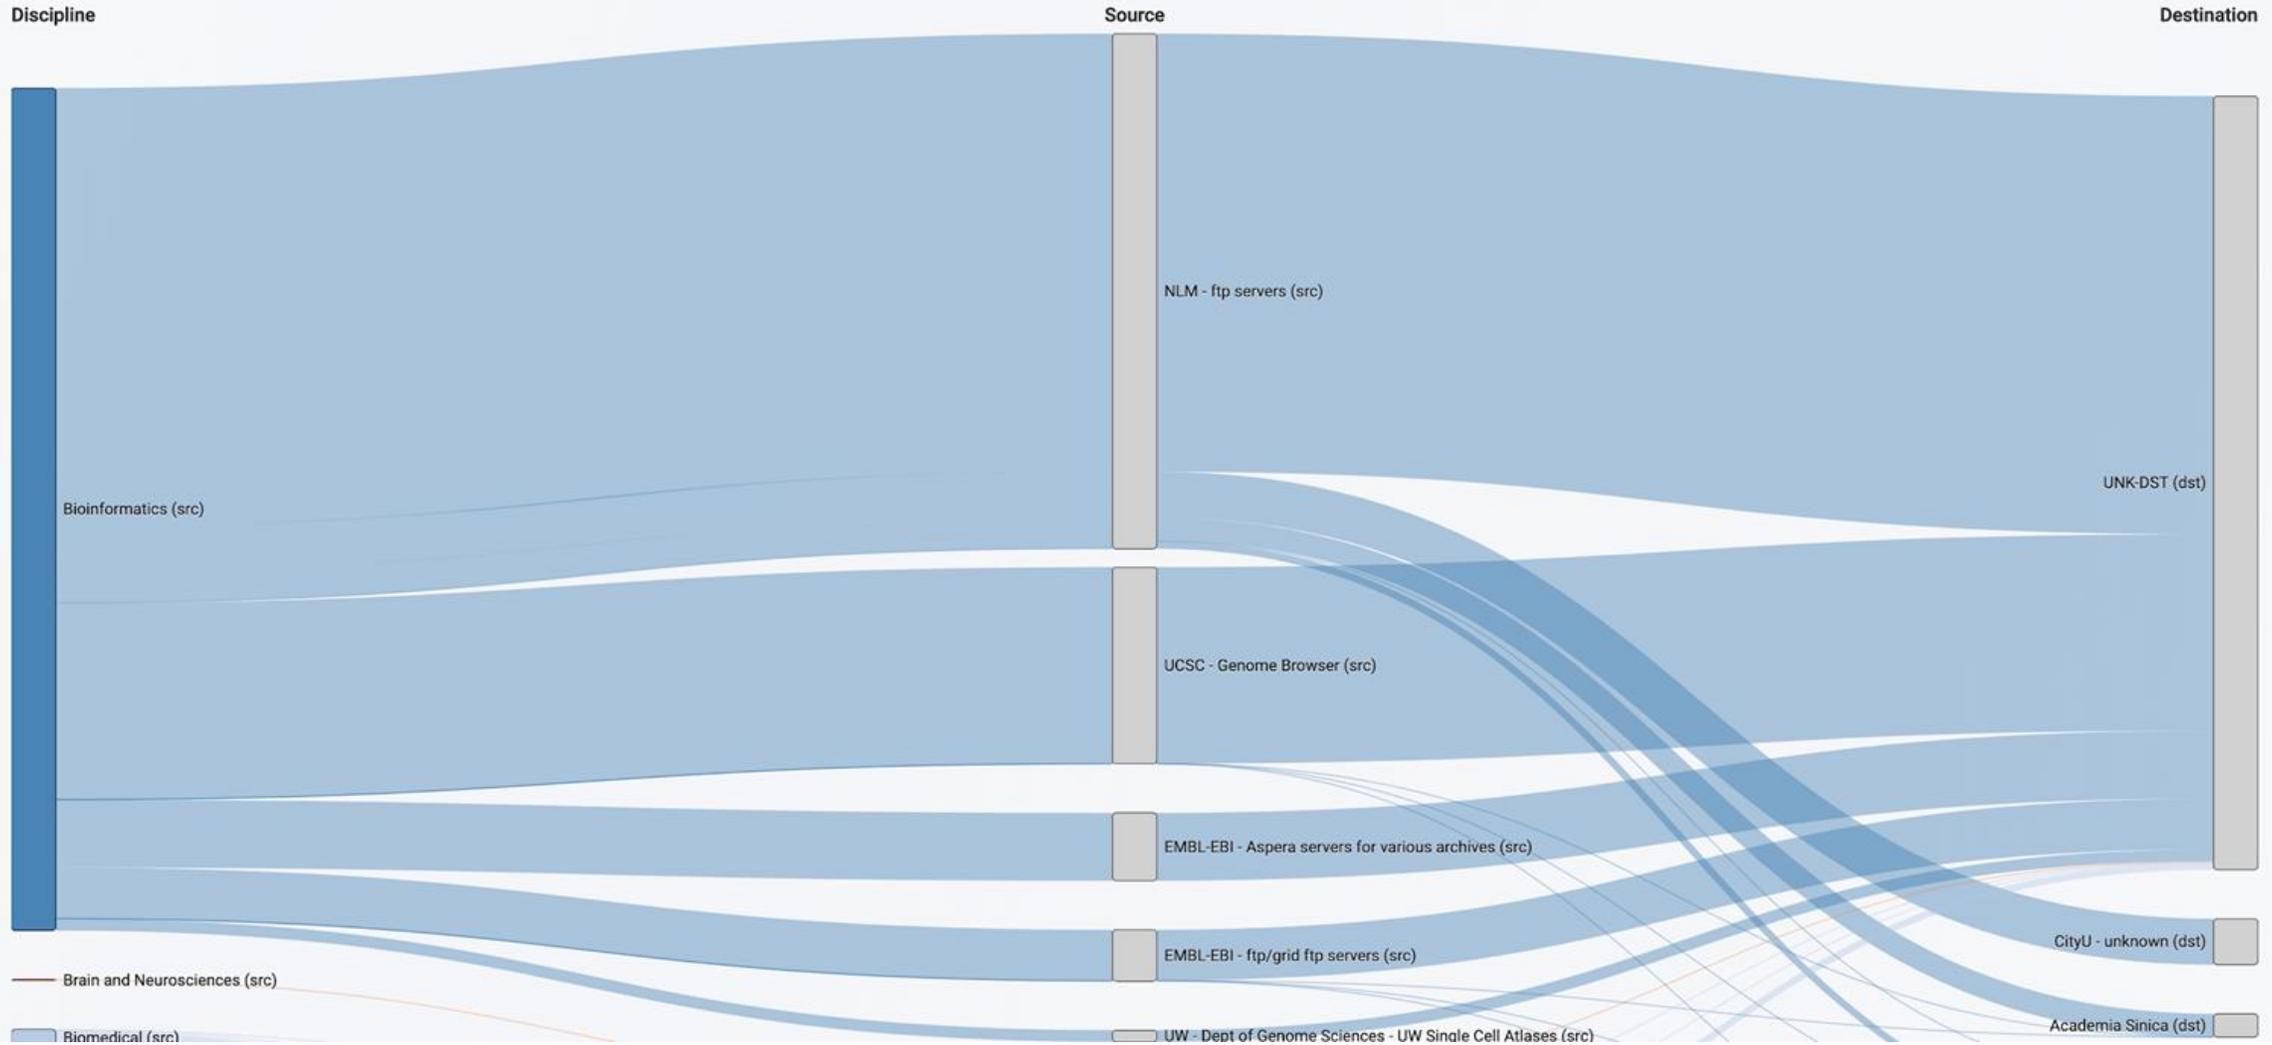


Bandwidth Patterns -



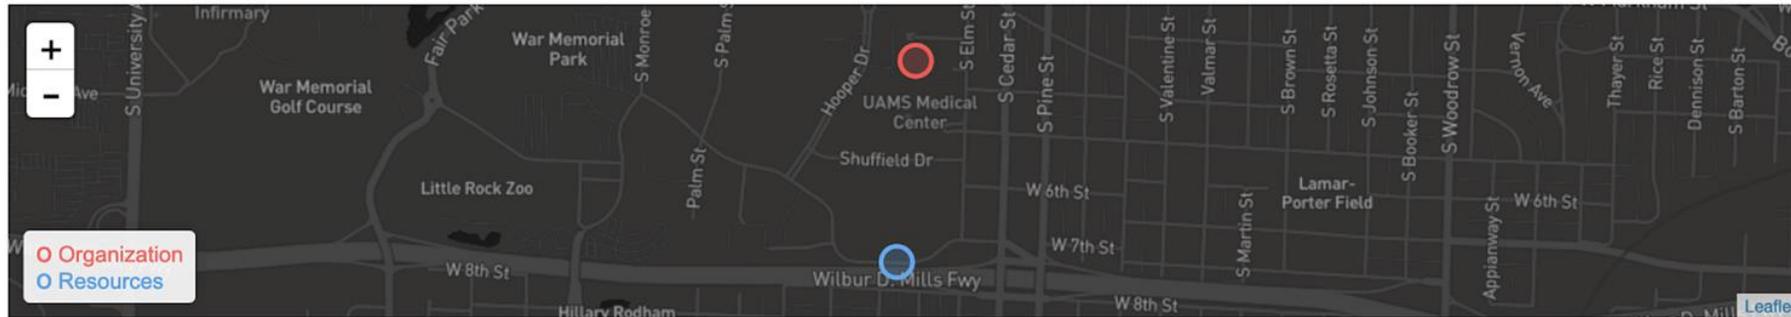
Flows by Science Discipline -

Last 30 days



University of Arkansas for Medical Sciences (UAMS)

The University of Arkansas for Medical Sciences (UAMS), Arkansas' only academic medical center, is part of the University of Arkansas System, a state-run university in the U.S. state of Arkansas. The main campus is located in Little Rock and consists of five colleges including one graduate school, seven institutes, a statewide network of community educational centers, and the UAMS Medical Center.



Organization Information

Details

Abbreviation UAMS
 Homepage <http://www.uams.edu/>
 Geolocation 34.749400, -92.320000
 Country United States
 Notes
 Contact
 Email



Resources		Events		
Resource Name	CIDR	Discipline	Role	Lat, Long
UAMS - The Cancer Imaging Archive (TCIA)	144.30.124.210/32	Biomedical	Storage	34.7449, -92.3205

[Add Resource](#)

TNRP Engagements and Deployments

- Nautilus “Hypercluster” deployment of perfSONAR mesh development and data collection
- Development environment for tools to enhance engagement opportunities
- Platform for systems to enhance research network resilience (detecting BGP hijacks) through ARTEMIS

Number of years:

1 5 10

1 2 3 4 5 6 7 8 9 10

State:

Missouri

.EDU site:

missouri.edu umkc.edu
mst.edu umsystem.edu
umsl.edu

Choose one:

interstate
 inter-institution
 all

Download

Dashboard

Raw data

Column visibility Copy Print Download Search:

link	award_id	title
Info	1339156	Polyploidy and Plasticity in the Crop Brassicas
Info	1351274	CAREER:Multi-Scale Meniscus Thin-Film Evaporation Enhancement using Hierarchical Tri-Porous Media
Info	1354609	Probing the Role of Glia in Neuronal Function and Behavior in Drosophila
Info	1355151	Collaborative Research:RCN:UrBioNet: A global network of urban biodiversity research and practice

Showing 1 to 4 of 227 entries

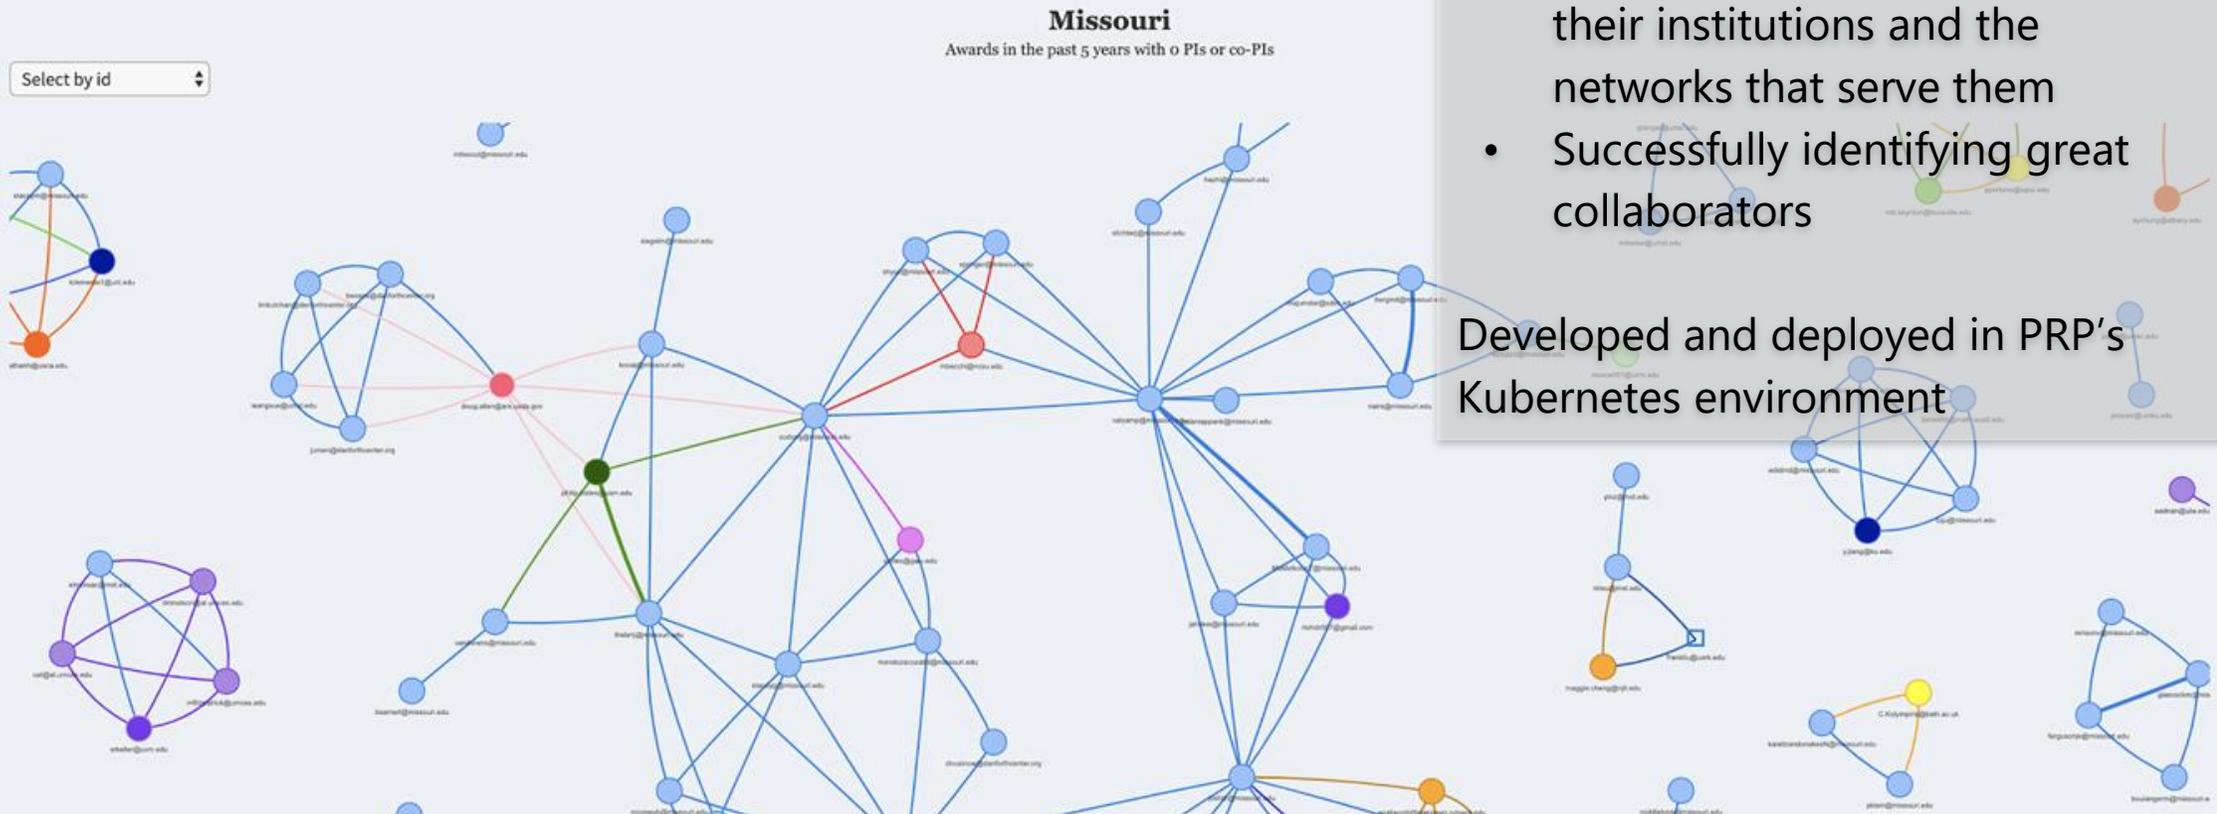


Social Network Analysis of Awards

Connecting network data such as ASNs to the institutions to help engage networkers

Interesting discoveries:

- Challenges of mapping co-PIs to their institutions and the networks that serve them
- Successfully identifying great collaborators



Developed and deployed in PRP's Kubernetes environment

ARTEMIS

self-managed BGP hijacking detection

<https://github.com/FORTH-ICS-INSPIRE/artemis>



BGP  STREAM

ARTEMIS



ARTEMIS Overview BGP Updates Hijacks Visualizations Admin Actions About Sign out

Hijacks Live Update:

Time Detected Status Profile Type Hijack AS # Peers Seen # ASes Infected Seen More

Time Detected	Status	Profile	Type	Hijack AS	# Peers Seen	# ASes Infected	Seen	More
2018-12-22 08:01	Warning	100-01-0076	1	1000	1	1	View	View
2018-12-22 08:02	Warning	100-01-0076	0	Unknown	1	0	View	View
2018-12-22 08:07	OK	100-01-200-024	0	8522	1	4	View	View
2018-12-22 08:08	OK	100-01-200-024	0	8899	1	4	View	View
2018-12-22 08:07	OK	100-01-0076	1	8408	1	0	View	View
2018-12-22 08:16	OK	100-01-200-024	0	8522	0	10	View	View
2018-12-22 08:16	OK	100-01-0076	0	8522	4	0	View	View

Showing 1 to 7 of 7 entries

Display / Refresh / Export / Under Migration / Withdraw / Outdated

Times are shown in your local time zone GMT-8 (America/Los_Angeles)

Sermpezis et al. , “ARTEMIS: Neutralizing BGP Hijacking within a Minute”
IEEE/ACM Transactions on Networking 2018



Center for Applied Internet Data Analysis
University of California San Diego

Foundation for Research and Technology-Hellas
University of Crete,



kubernetes

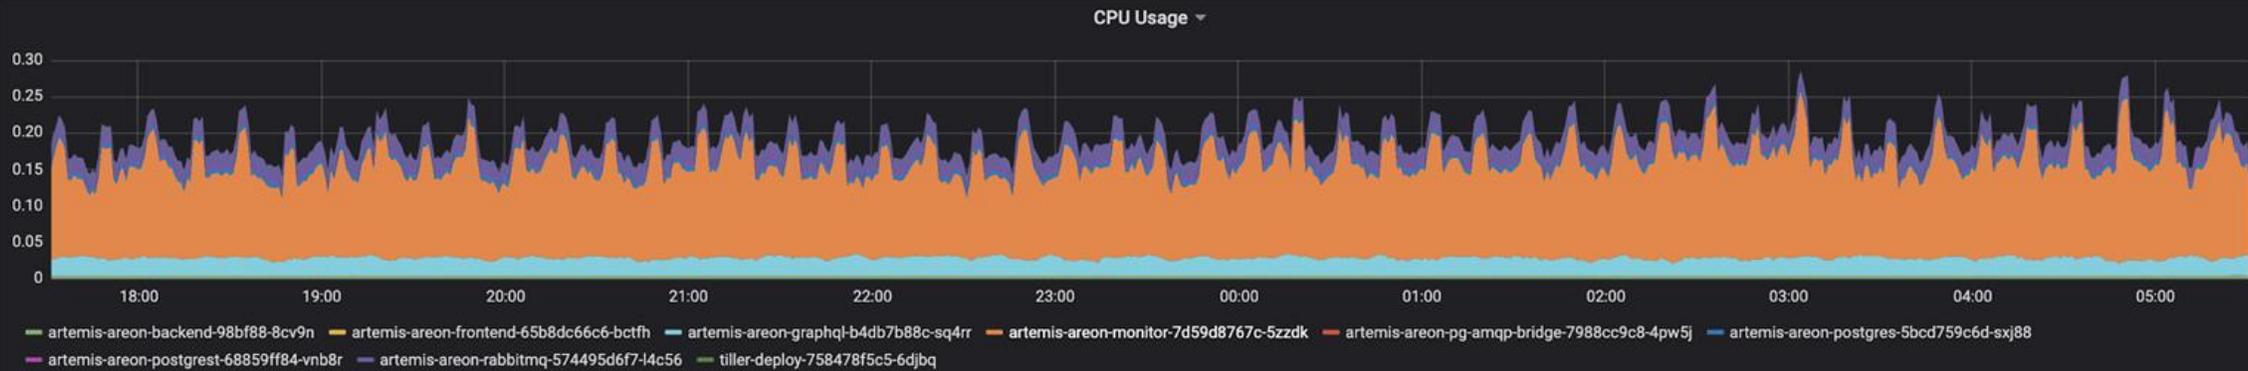
PRP PACIFIC RESEARCH
PLATFORM



NSF OAC-1848641 — Sep 2018 - Aug 2019
**Experimental Deployment of the ARTEMIS
BGP Hijacking Detection Prototype in
Research and Educational Networks**



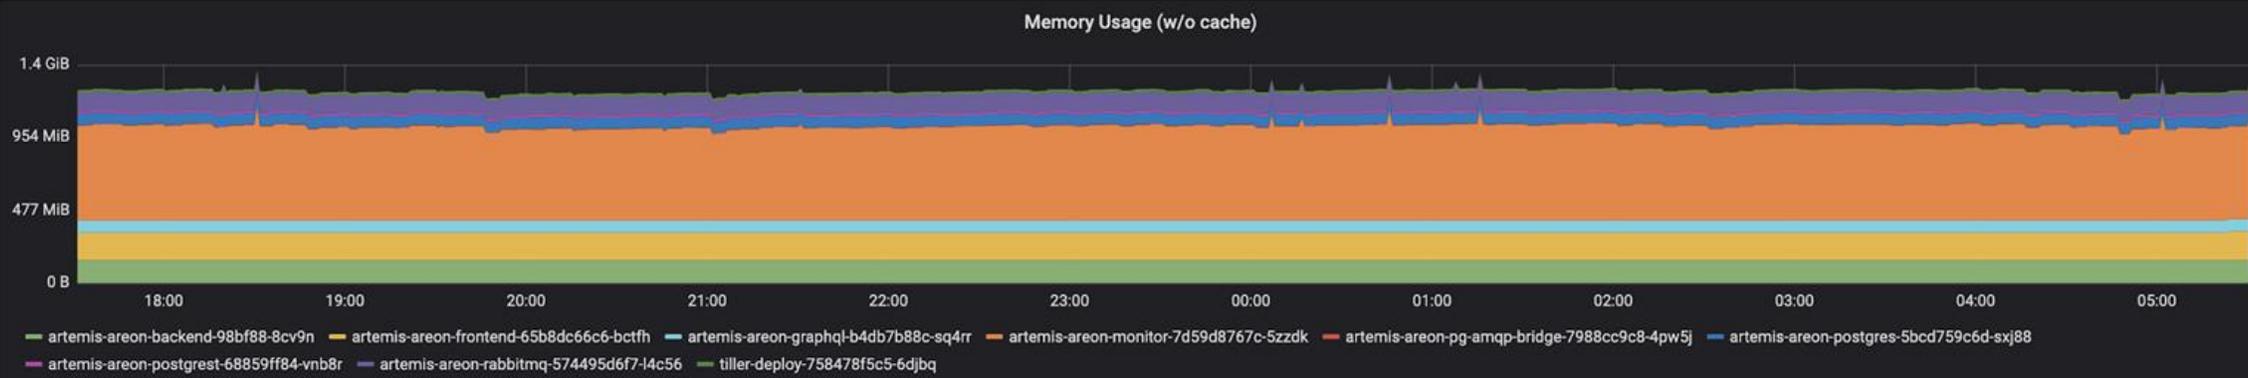
CPU Usage



CPU Quota

Pod	CPU Usage
artemis-areon-rabbitmq-574495d6f7-l4c56	0.02
artemis-areon-postgres-5bcd759c6d-sxj88	0.01
artemis-areon-pg-amqp-bridge-7988cc9c8-4pw5j	0.00
artemis-areon-graphql-b4db7b88c-sq4rr	0.03
artemis-areon-frontend-65b8dc66c6-bctfh	0.00
artemis-areon-backend-98bf88-8cv9n	0.00

Memory Usage



Summary

- **Gaining insight into patterns and projects utilizing the connections**
- **Learning better ways to engage and discover scientific workflows**
- **Encouraging conversations that have led to much greater amounts of network metric sharing across multiple regionals**
- **Simplifying the tougher aspects of research network measurement and awareness**