# PA Science DMZ Overview
# CyberAccelerate Workshop

Ken Miller, Chief Technology Officer

KeystoneREN

# Science DMZ:
# A Scalable Network Design Pattern for Optimizing Science Data Transfers

**Science DMZ Overview:**

- Network architecture optimized for scientific data transfers or specific network data architecture requirements
- Positioned at or near the campus/laboratory's network perimeter
- Typically tailored for high-performance science applications (e.g., data transfer, remote experiments)

**Development & Purpose:**

- Created by ESnet engineers
- Addresses common data transfer performance issues in research institutions
- Optimized for high-volume data transfers, low latency experiment control, and real-time data visualization

**Key Features:**

- Scalable and incrementally deployable
- Flexible design for various data security compliance requirements
- Adaptable to advanced technologies like 400 Gigabit Ethernet, virtual circuits, network overlays, security enclaves, etc.

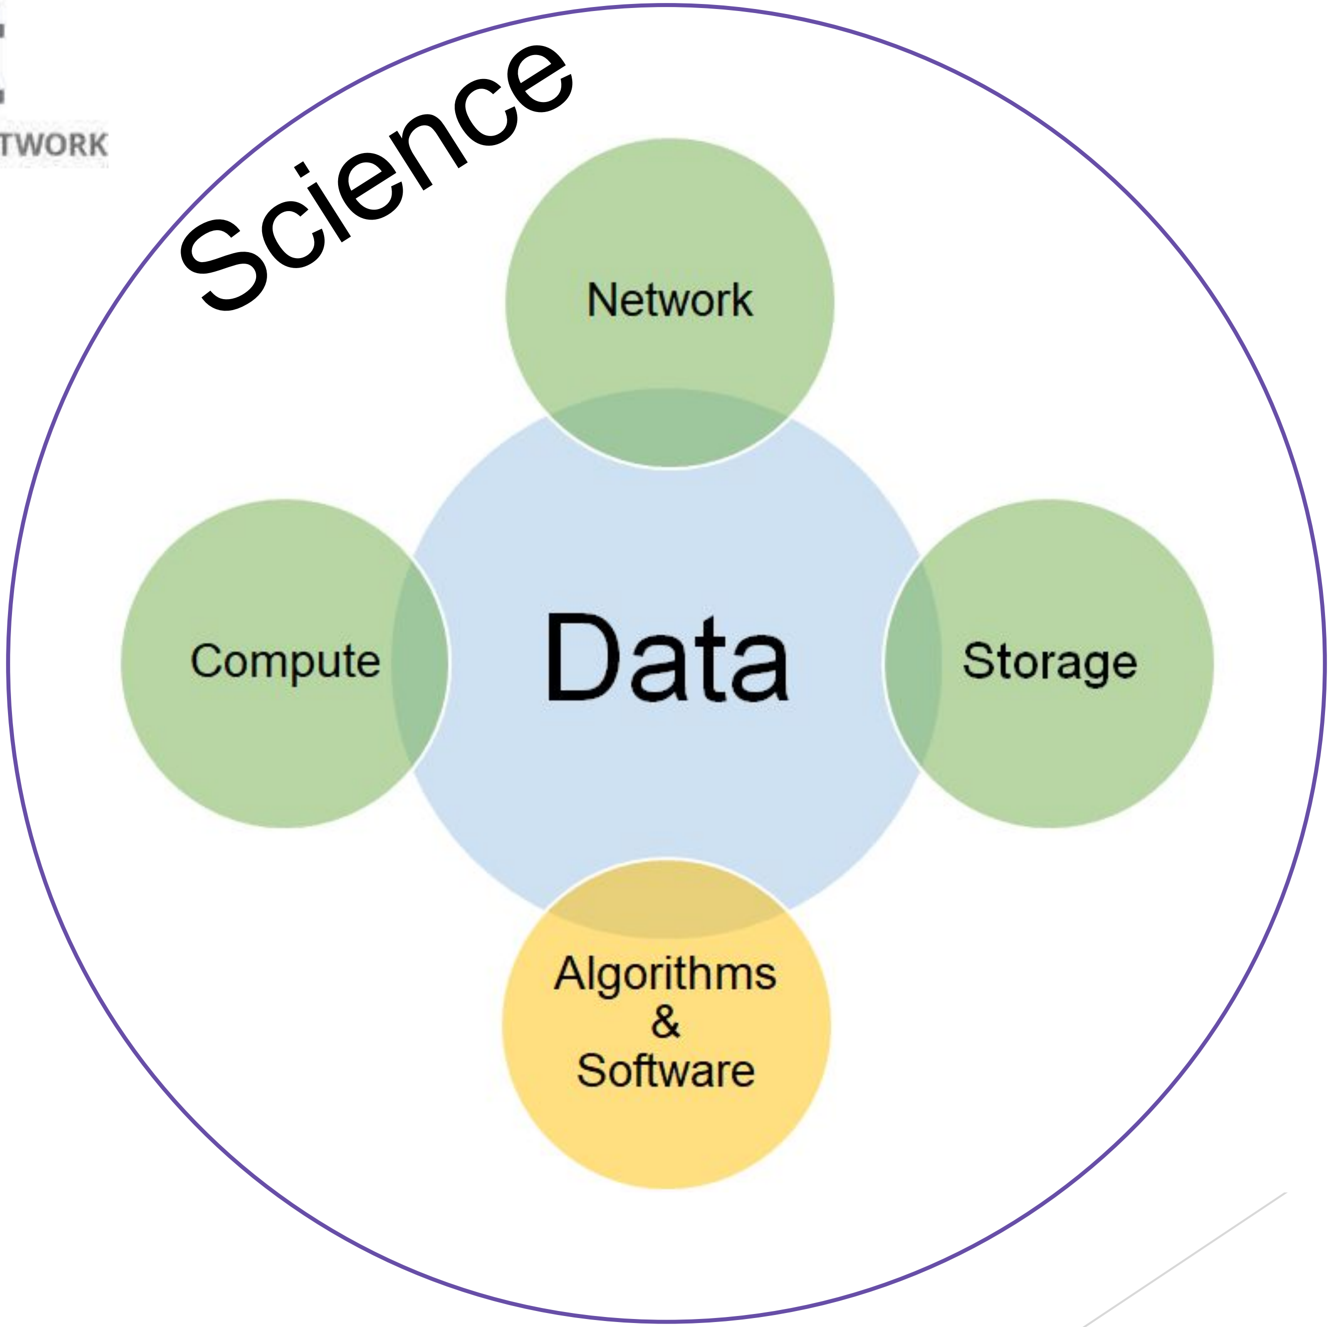
Science
Network
Compute
Data
Storage
Algorithms
&
Software

# Th

Consi

- “Fr
  - 
  - 
  - 
  - 
- Dec
  - 
  - 
- Per
  - 
- Eng

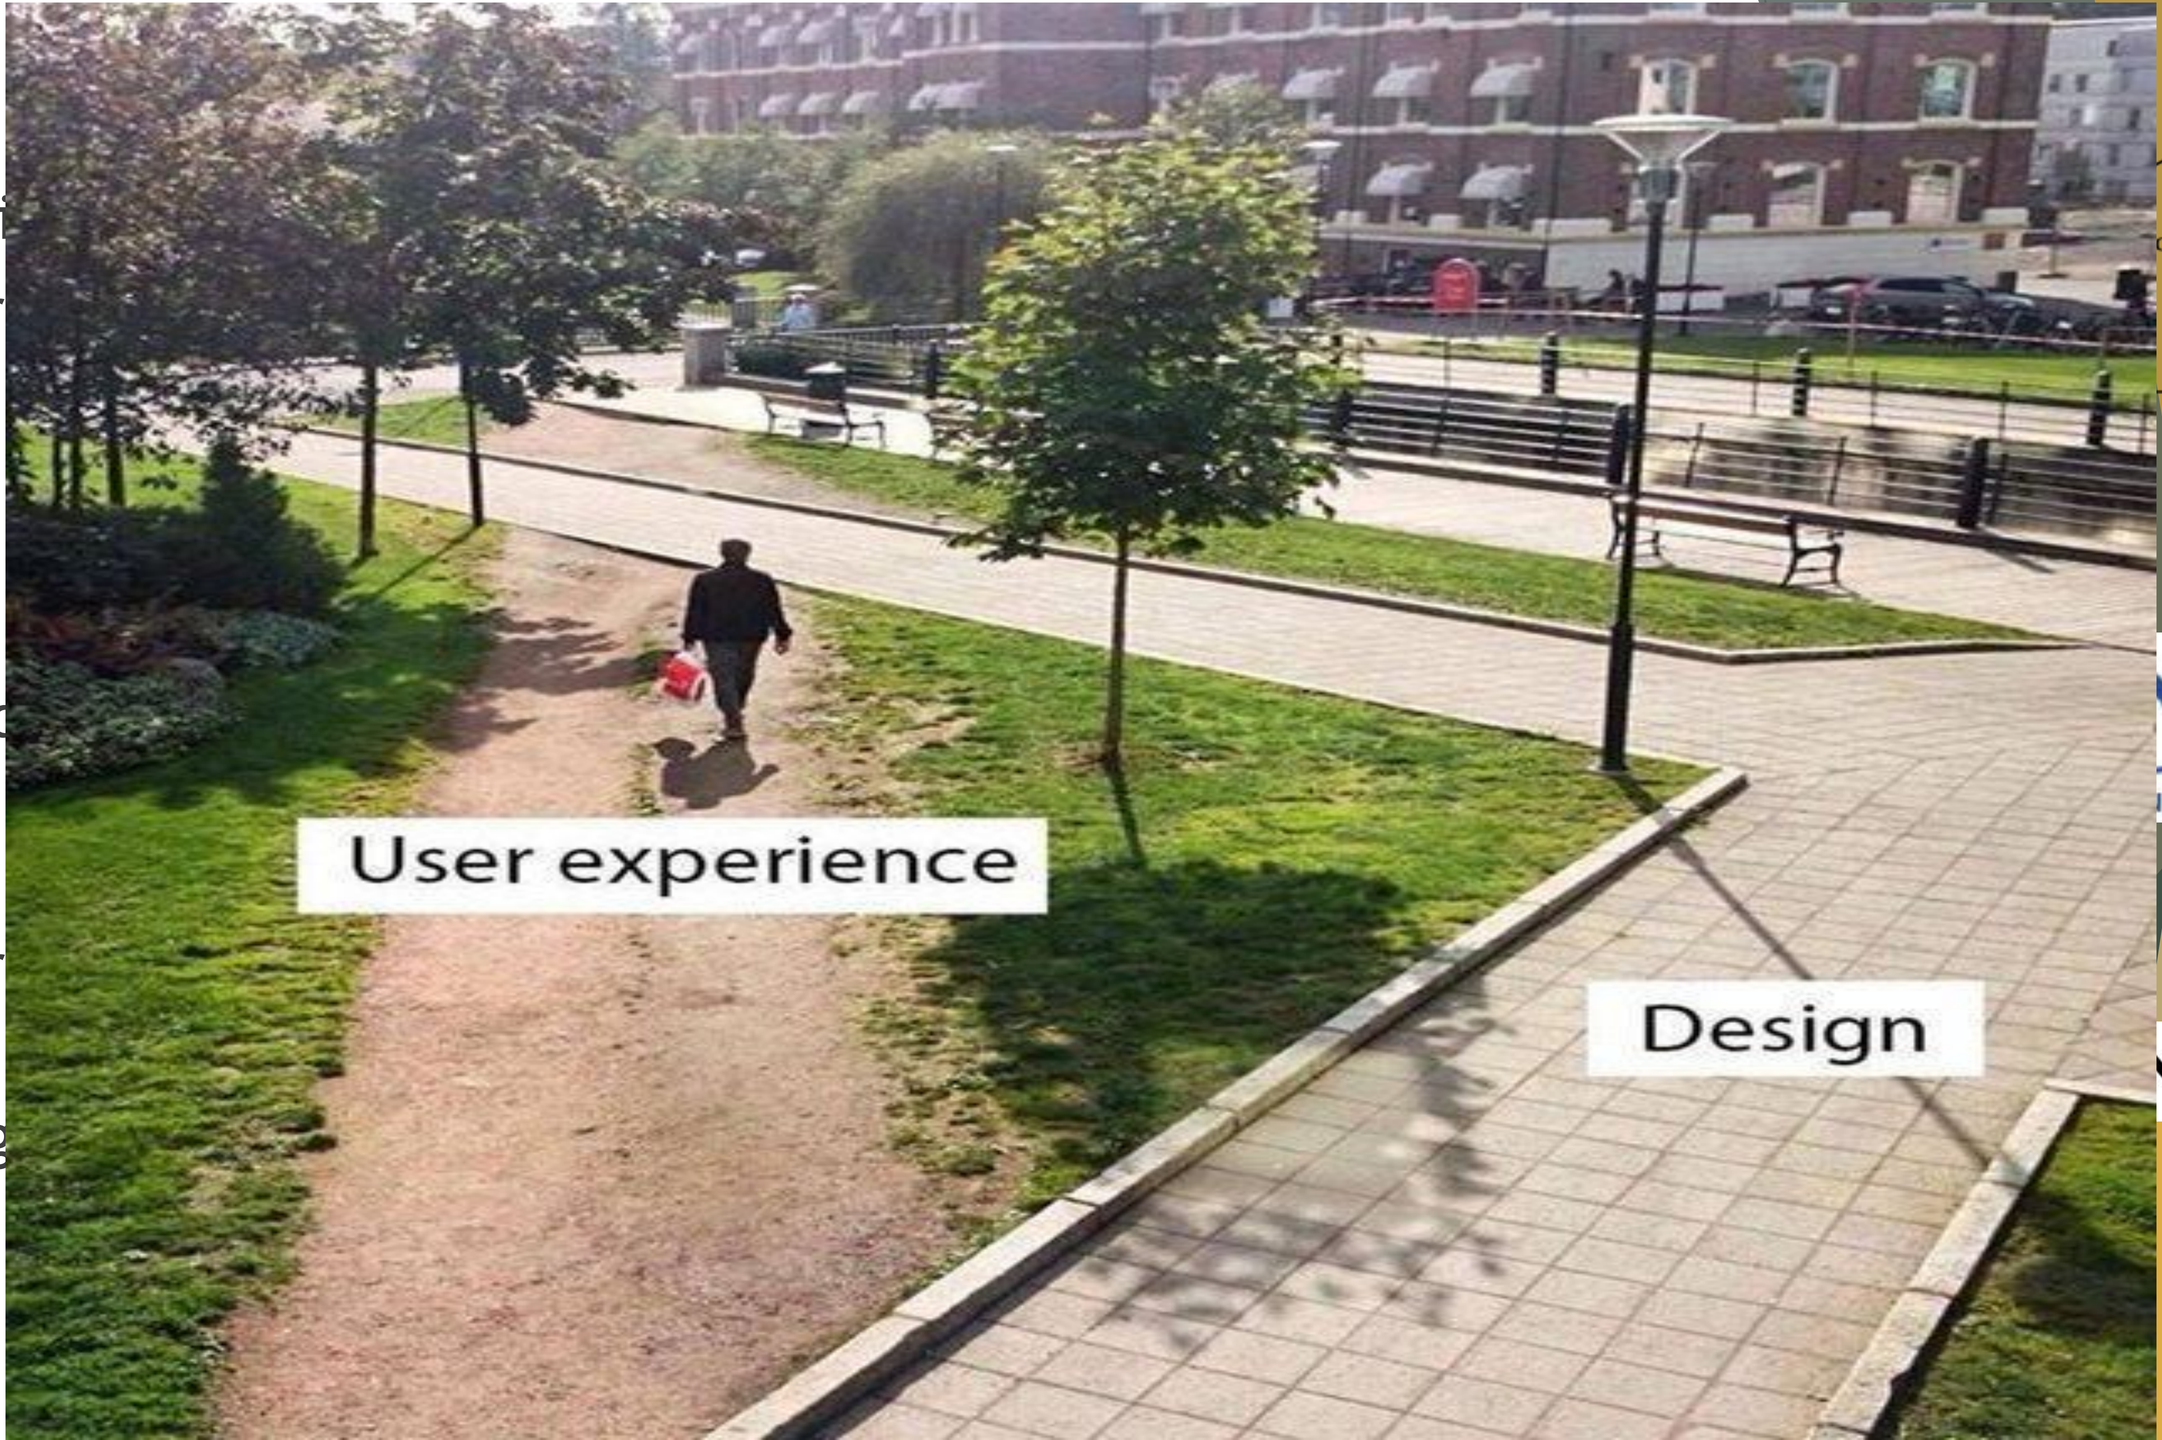

NAR

# A small amount of packet loss makes a huge difference in TCP performance

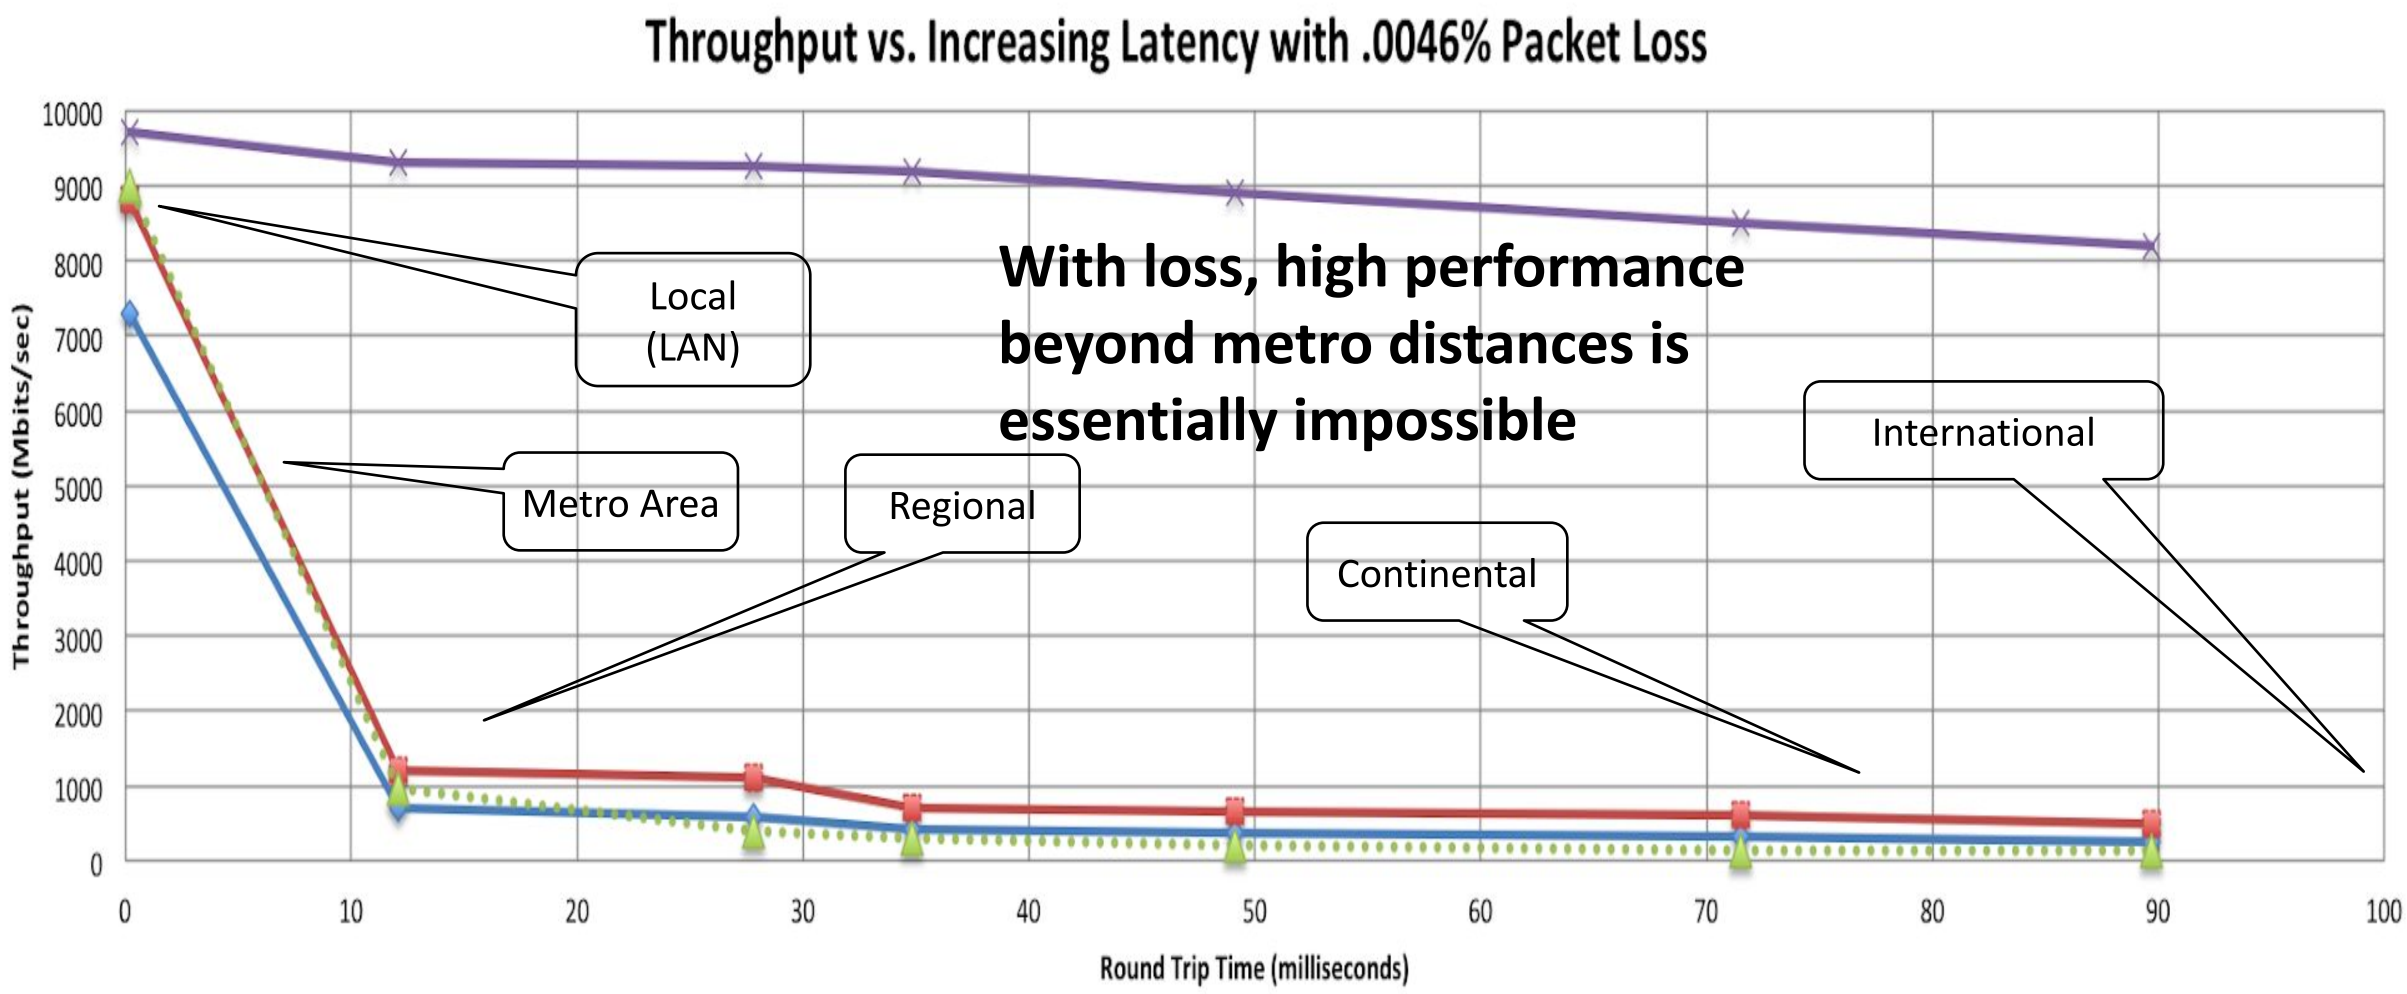

**With loss, high performance beyond metro distances is essentially impossible**

**Measured (TCP Reno)** **Measured (HTCP)** **Theoretical (TCP Reno)** **Measured (no loss)**

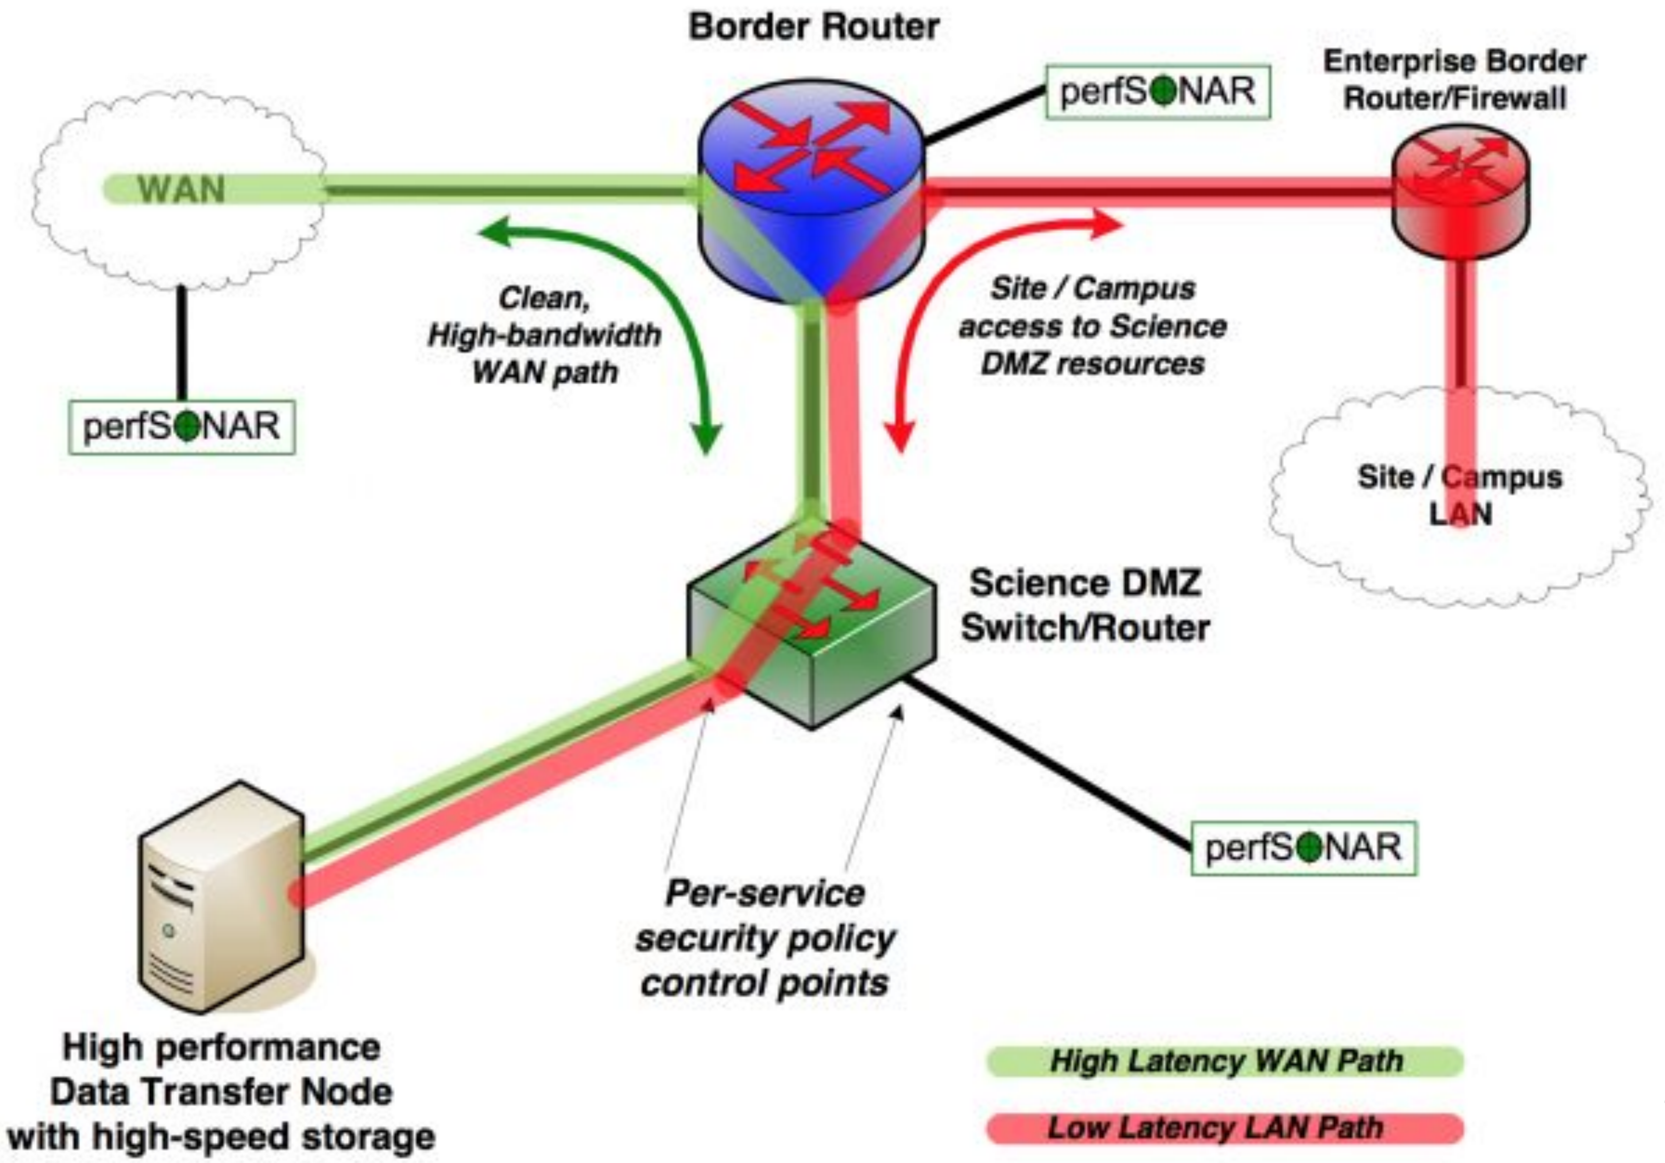
Border Router
perfSONAR
Enterprise Border Router/Firewall
WAN
perfSONAR
Clean, High-bandwidth WAN path
Site / Campus access to Science DMZ resources
Site / Campus LAN
Science DMZ Switch/Router
perfSONAR
Per-service security policy control points
High performance Data Transfer Node with high-speed storage
High Latency WAN Path
Low Latency LAN Path

Simple Science DMZ

Science DMZ is the HOV lane for research data and workflows

# PA-Science DMZ PROJECT OVERVIEW

- Frictionless Science DMZ Network Paths
  - Goal - establish the foundation for a statewide Pennsylvania Regional Science DMZ (PA-DMZ) that enables and enhances access for under resourced PA institutions of higher education to cyberinfrastructure-based resources and services in support of science driven research and education applications.
- Grant supports
  - Networking hardware and connectivity
  - Installation and support for 2+ years (organizations to provide support years 3-5)
  - Broader Impacts and Research Enablement

# 2023 AWARD HIGHLIGHTS

- $1.1M funding - NSF Award #2346589
- 5 partners
- PA Science DMZ for Under-resourced Institutions
  - Existing 1-2Gb/s Internet only
  - Adding 10/25Gb/s router, 10Gb/s Internet2, with 10G perfSONAR and 10G DTNs
- Install and Operational in 2024
- Research Enablement in 2024/2025
- Growth and Expansion in 2025

# PA SCIENCE DMZ 2023 GRANT

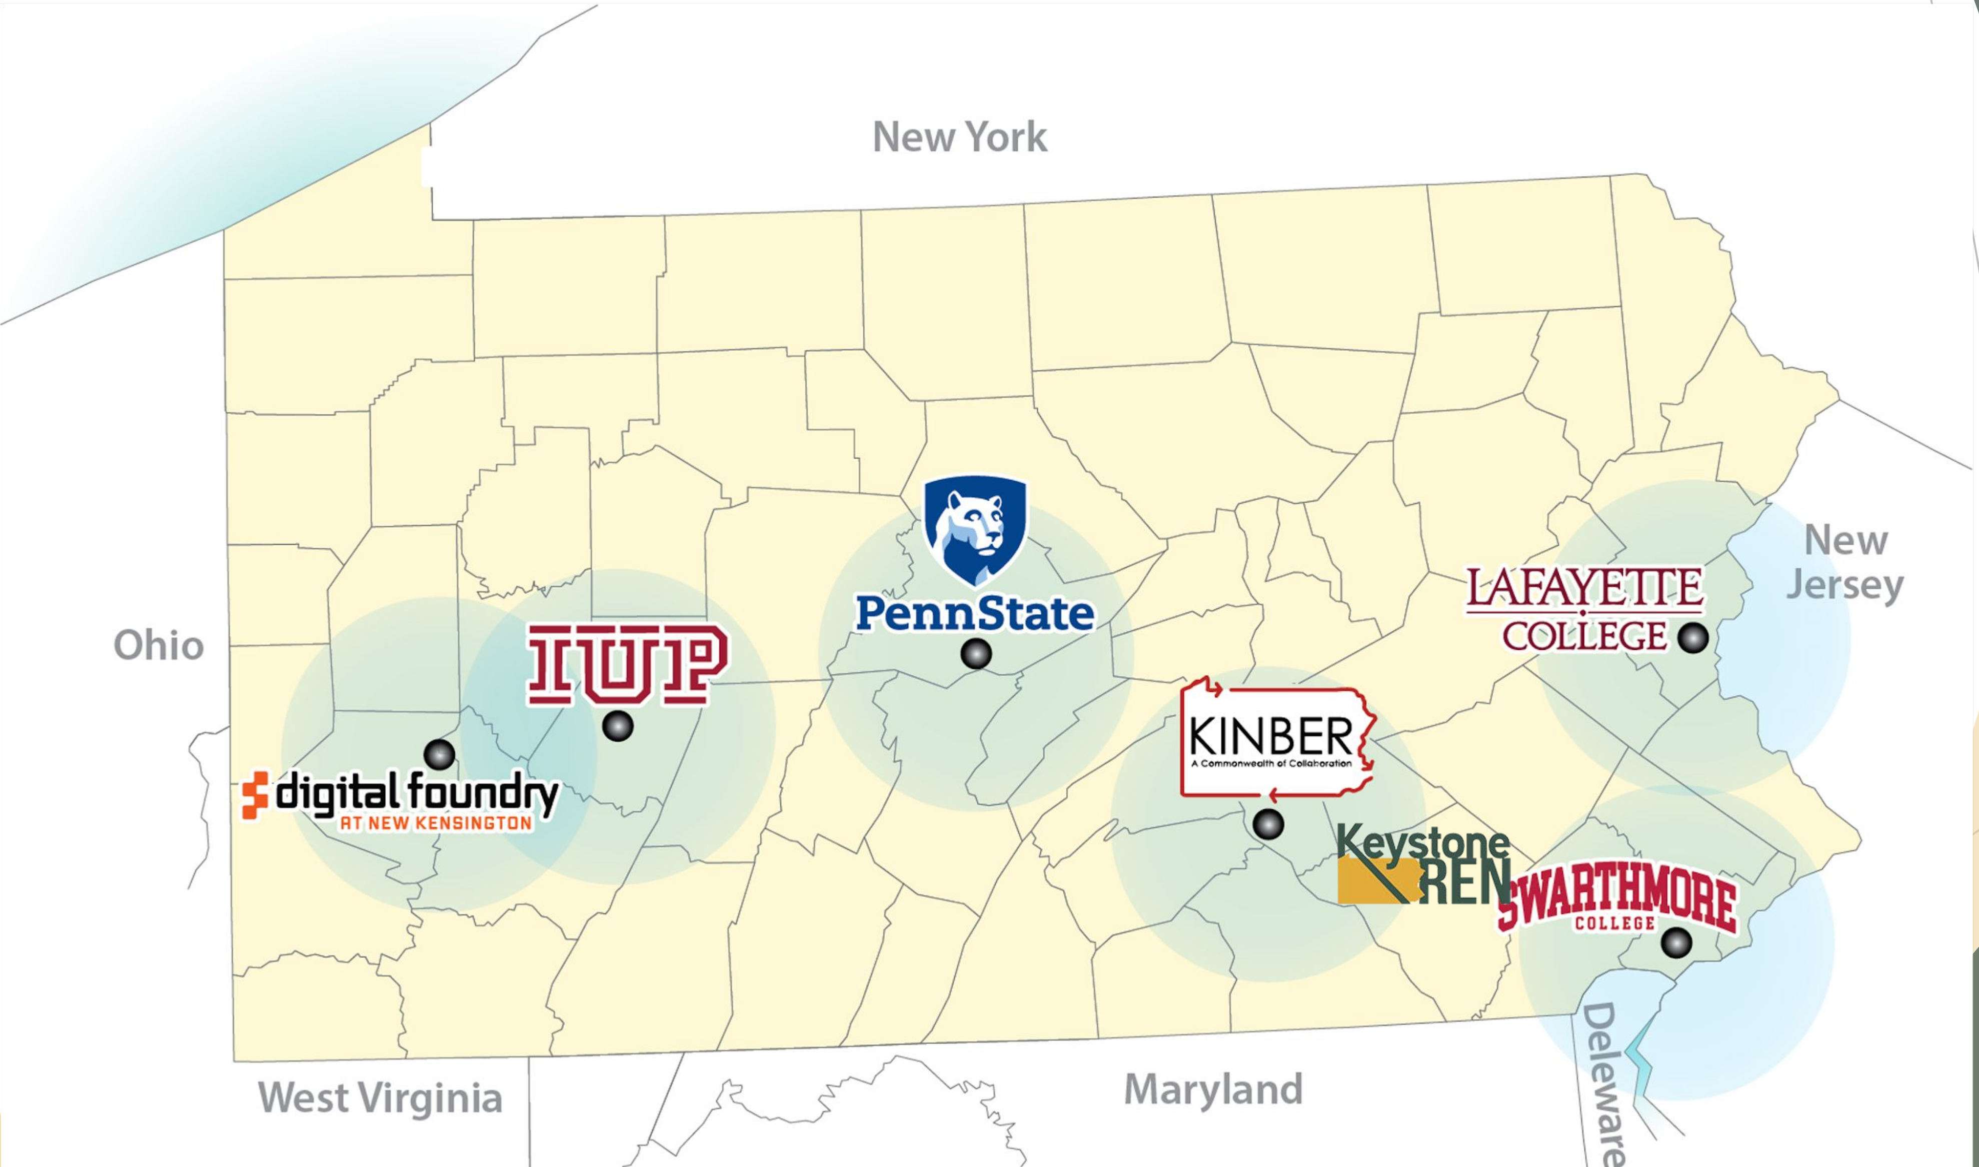

# PA SCIENCE DMZ 2024 GRANT Proposal

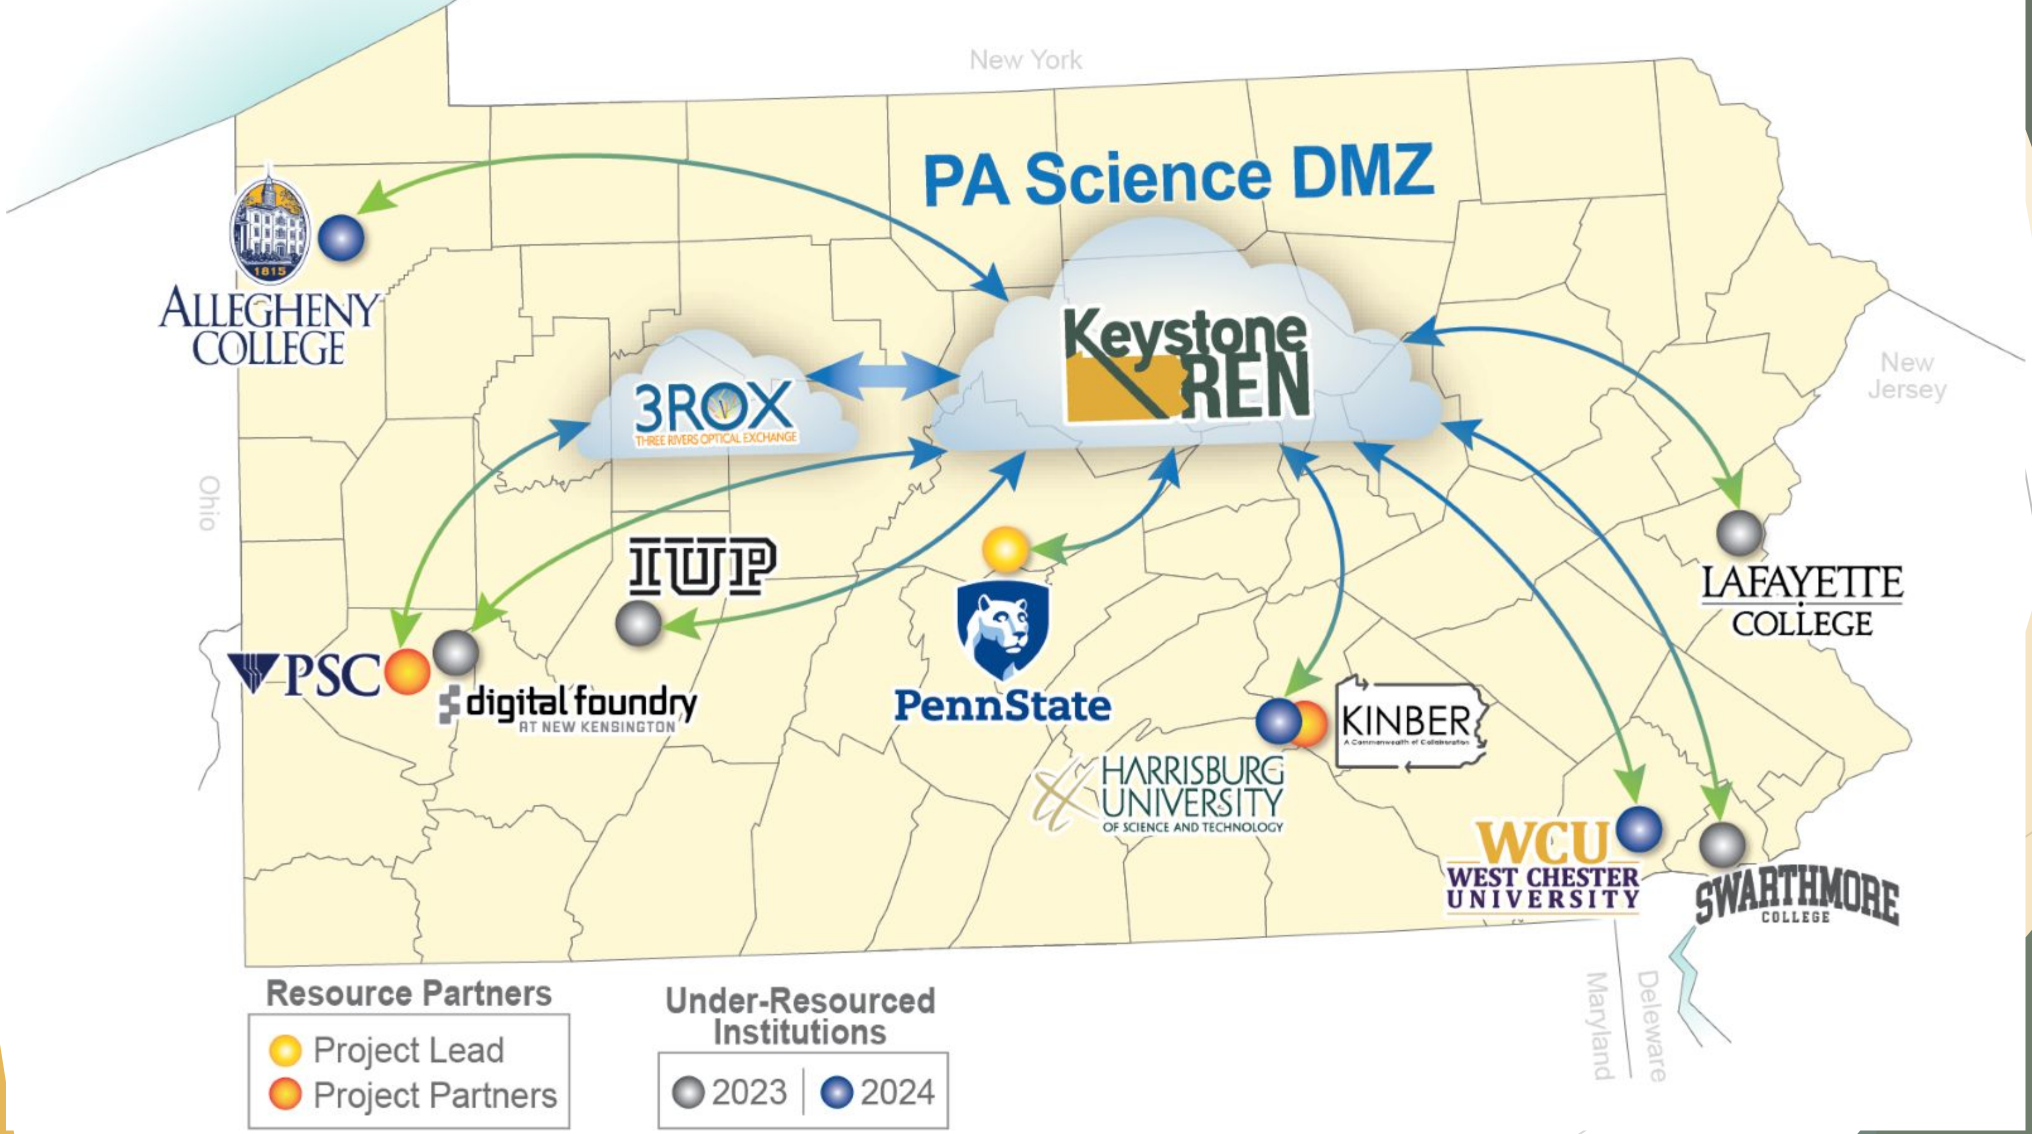

# PENNSYLVANIA'S POTENTIAL SCIENCE DMZ SITES

In PA – There have been 24 CC* awards (as of Jan 2024)

Many under-resourced institutions

Chance to provide researc enabling services to scale discovery

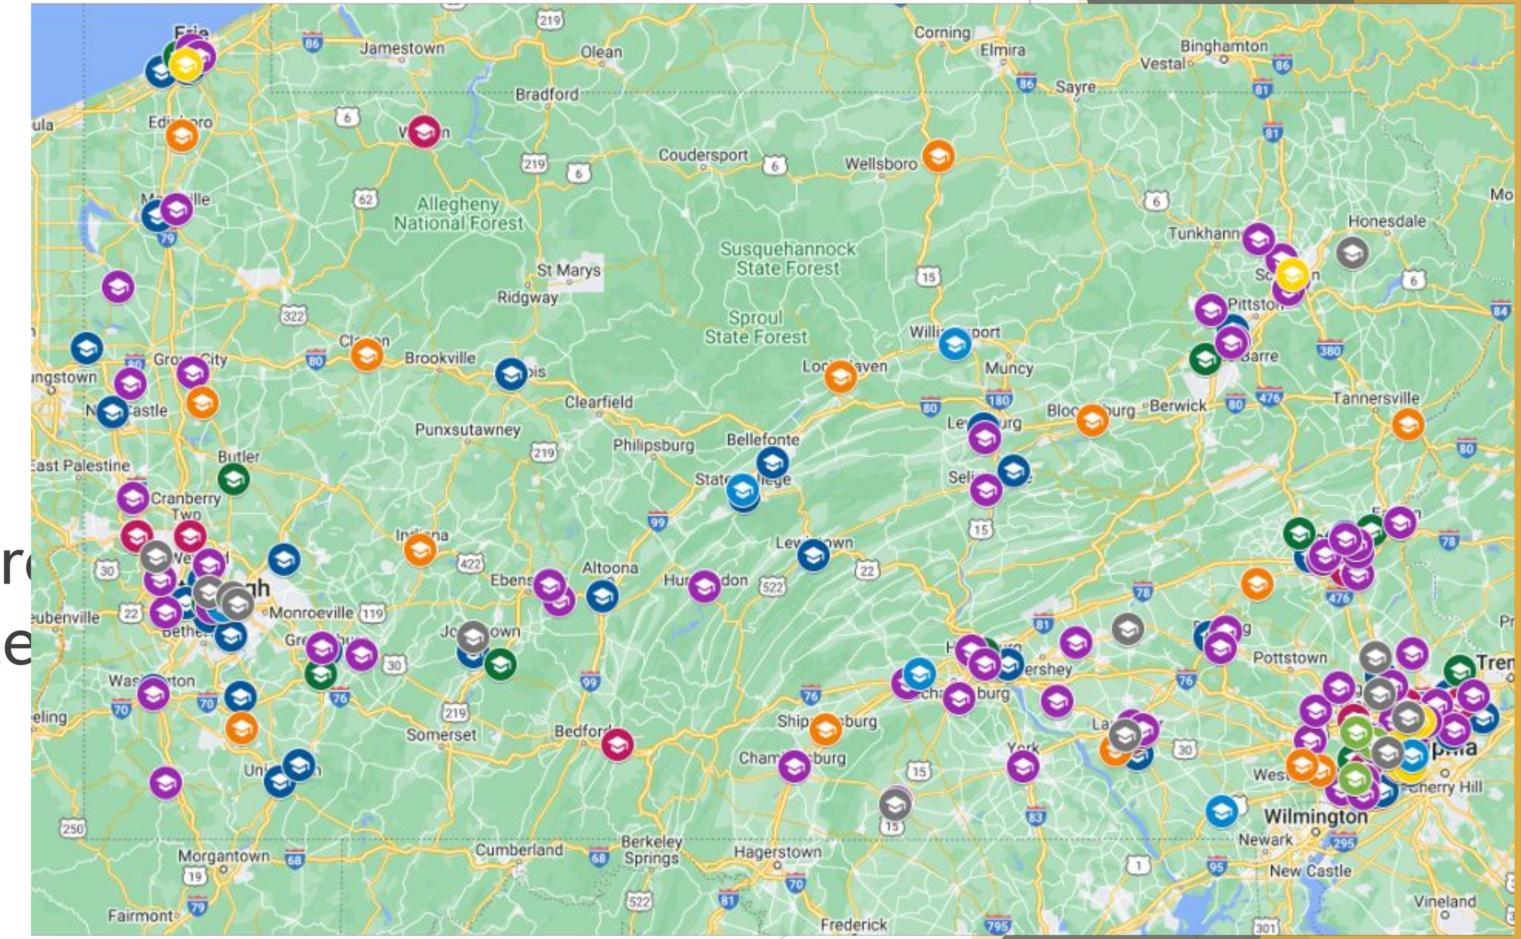

# CYBERSECURITRY ON PA SCIENCE DMZ

- Implicit Deny all/block all traffic, ACLs for IPv4 and IPv6
- ACLs opened as science drivers are identified and documented
- ACL accounting on all accepted and denied packets logged to campus security
- All accepted packets mirrored to campus security
- sFlow or Netflow/IPFIX will be captured on PA Science DMZ equipment
- Routing Optimization to prefer R&E networks only

# PA SCIENCE DMZ PERFORMANCE

- perfSONAR testing IPv4 and IPv6
  - MTU 9000 verification or at least MTU consistency
  - Throughput = iperf3 (single and multi threaded) to verify network capacity
  - Latency = Owe-way and round trip
  - Traceroute to make sure traffic is on R&E paths only
- Data Transfer Node testing
  - Once network performance is validation, DTN will be tested with datasets toto well tuned endpoints at ESnet measure against Data Transfer Scorecard – 1-3 TB/hr or 2-6 Gb/s
  - Utilize the Modern Research Data Portal with Globus and ESnet's data architecture design pattern.  Free Code here
  - Collaborate with Science Driver to validate data transfer against Data Transfer Scorecard

# DATA TRANSFER SCORECARD

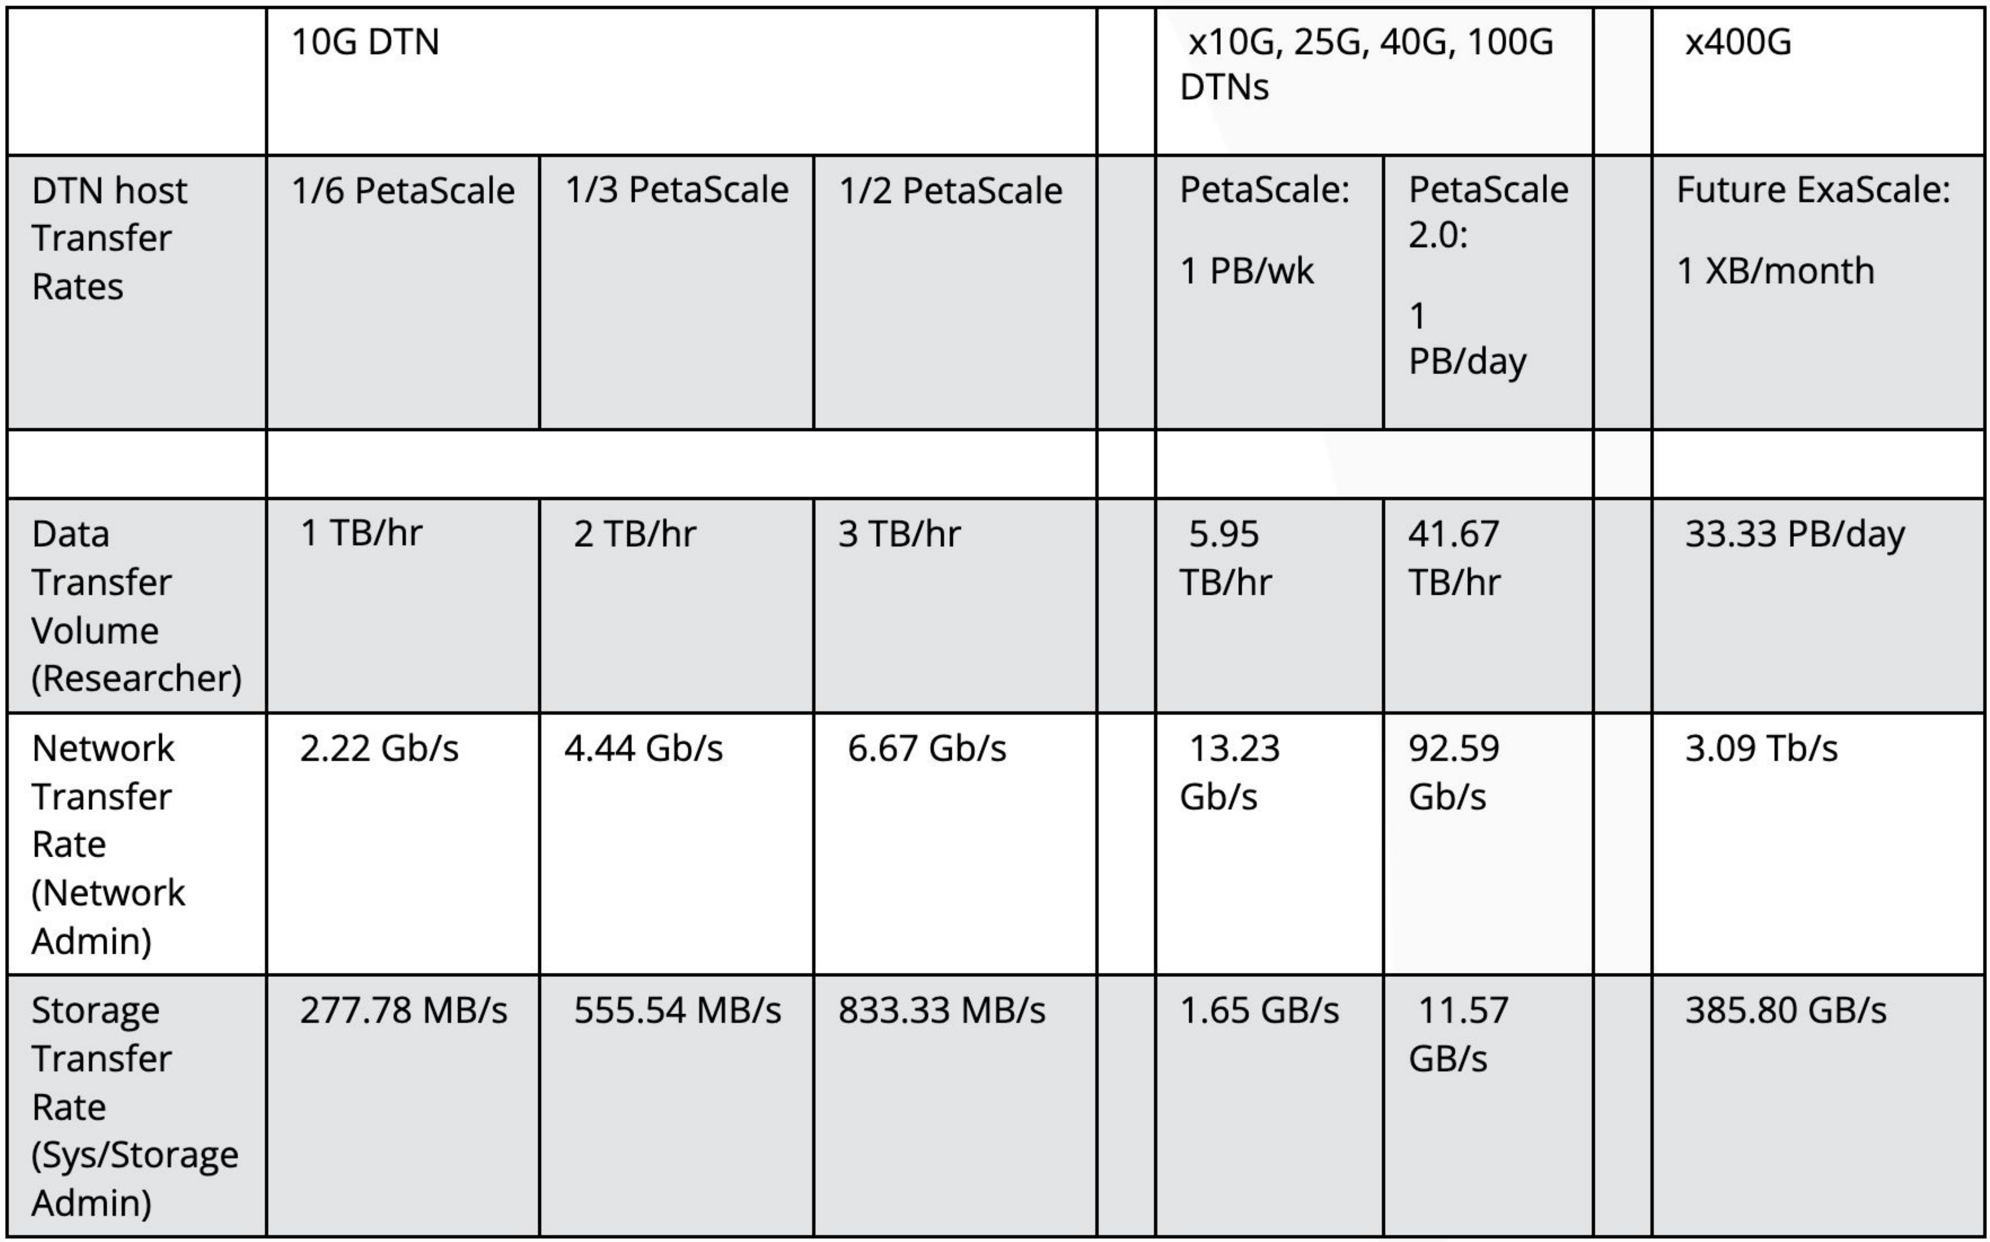

| | 10G DTN | | | | x10G, 25G, 40G, 100G DTNs | | | x400G |
|---|---|---|---|---|---|---|---|---|
| DTN host Transfer Rates | 1/6 PetaScale | 1/3 PetaScale | 1/2 PetaScale | | PetaScale:<br>1 PB/wk | PetaScale 2.0:<br>1 PB/day | | Future ExaScale:<br>1 XB/month |
| | | | | | | | | |
| Data Transfer Volume (Researcher) | 1 TB/hr | 2 TB/hr | 3 TB/hr | | 5.95 TB/hr | 41.67 TB/hr | | 33.33 PB/day |
| Network Transfer Rate (Network Admin) | 2.22 Gb/s | 4.44 Gb/s | 6.67 Gb/s | | 13.23 Gb/s | 92.59 Gb/s | | 3.09 Tb/s |
| Storage Transfer Rate (Sys/Storage Admin) | 277.78 MB/s | 555.54 MB/s | 833.33 MB/s | | 1.65 GB/s | 11.57 GB/s | | 385.80 GB/s |

# SCIENCE DRIVER METRICS & OUTCOMES

- Baseline: Gather existing data transfer bottleneck or limitations
- Top Source/Destination
  - IPs/Collaborators
  - ASNs/Sites
  - Applications
- Total Science Data Transferred
- How has Science Improved?
- Develop a performant data architecture to assist others within PA

# How to use the PA DMZ

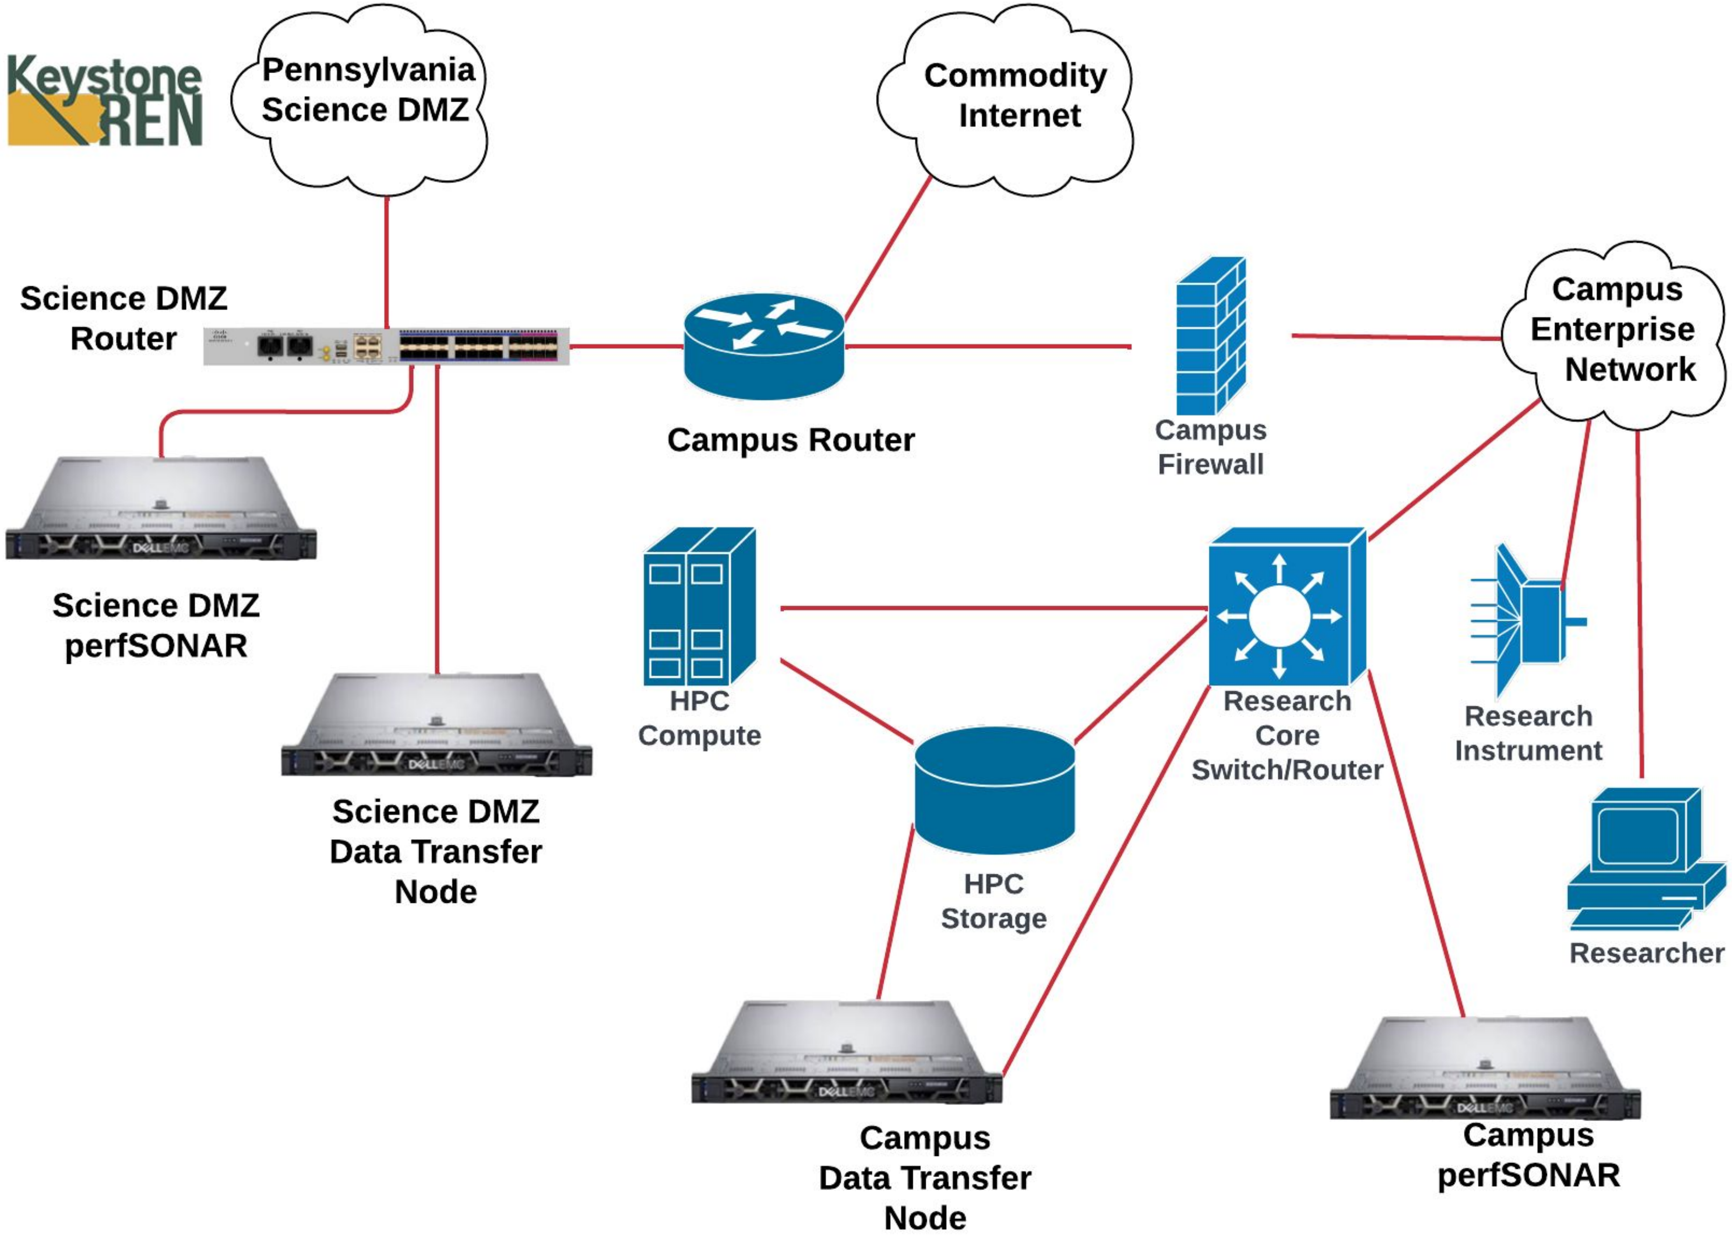

# Campus Science DMZ and Enterprise Network

# PA Science DMZ connectivity and usage

- Research and Education Networks are built for secure performance faster data transport, lowest latency possible, most direct paths possible, and jumbo network data packets
  - Enterprise networks are build for general connectivity
  - Data Center networks are build for short bursts of traffic within close buildings
- You use the PA Science DMZ by
  - Leveraging equipment connected to it
  - Prioritizing campus traffic over Research and Education networks instead of commodity
- The PA Science DMZ has dedicated Data Transfer Nodes connected to the network which are tuned for high performance data transfers which run Globus Connect Server Software
- We are capable of hosting research equipment in our data centers and connecting directly to the PA Science DMZ

# Globus Connect Server (GCS) is highly performant

- Parallel Data Transfer Streams
  - GCS leverages multiple parallel data streams, which allows for faster transfers compared to single-threaded or traditional file transfer methods. This parallelism maximizes throughput over networks by distributing the data transfer load across multiple channels simultaneously.
- Optimized for High-Latency Networks
  - GCS is designed to mitigate the effects of latency by using advanced techniques like pipelining and tuning buffer sizes, ensuring that transfers remain efficient even on high-latency networks.
- Automatic Fault Recovery
- Built-in Data Integrity Verification
- Integration with High-Performance Network Infrastructures Globally
- Efficient Use of DTN Resources
- Simple User Interface & Automation of workflows
- Security without Compromise
- Efficient Large File Handling

# Globus Intuitive web application interface

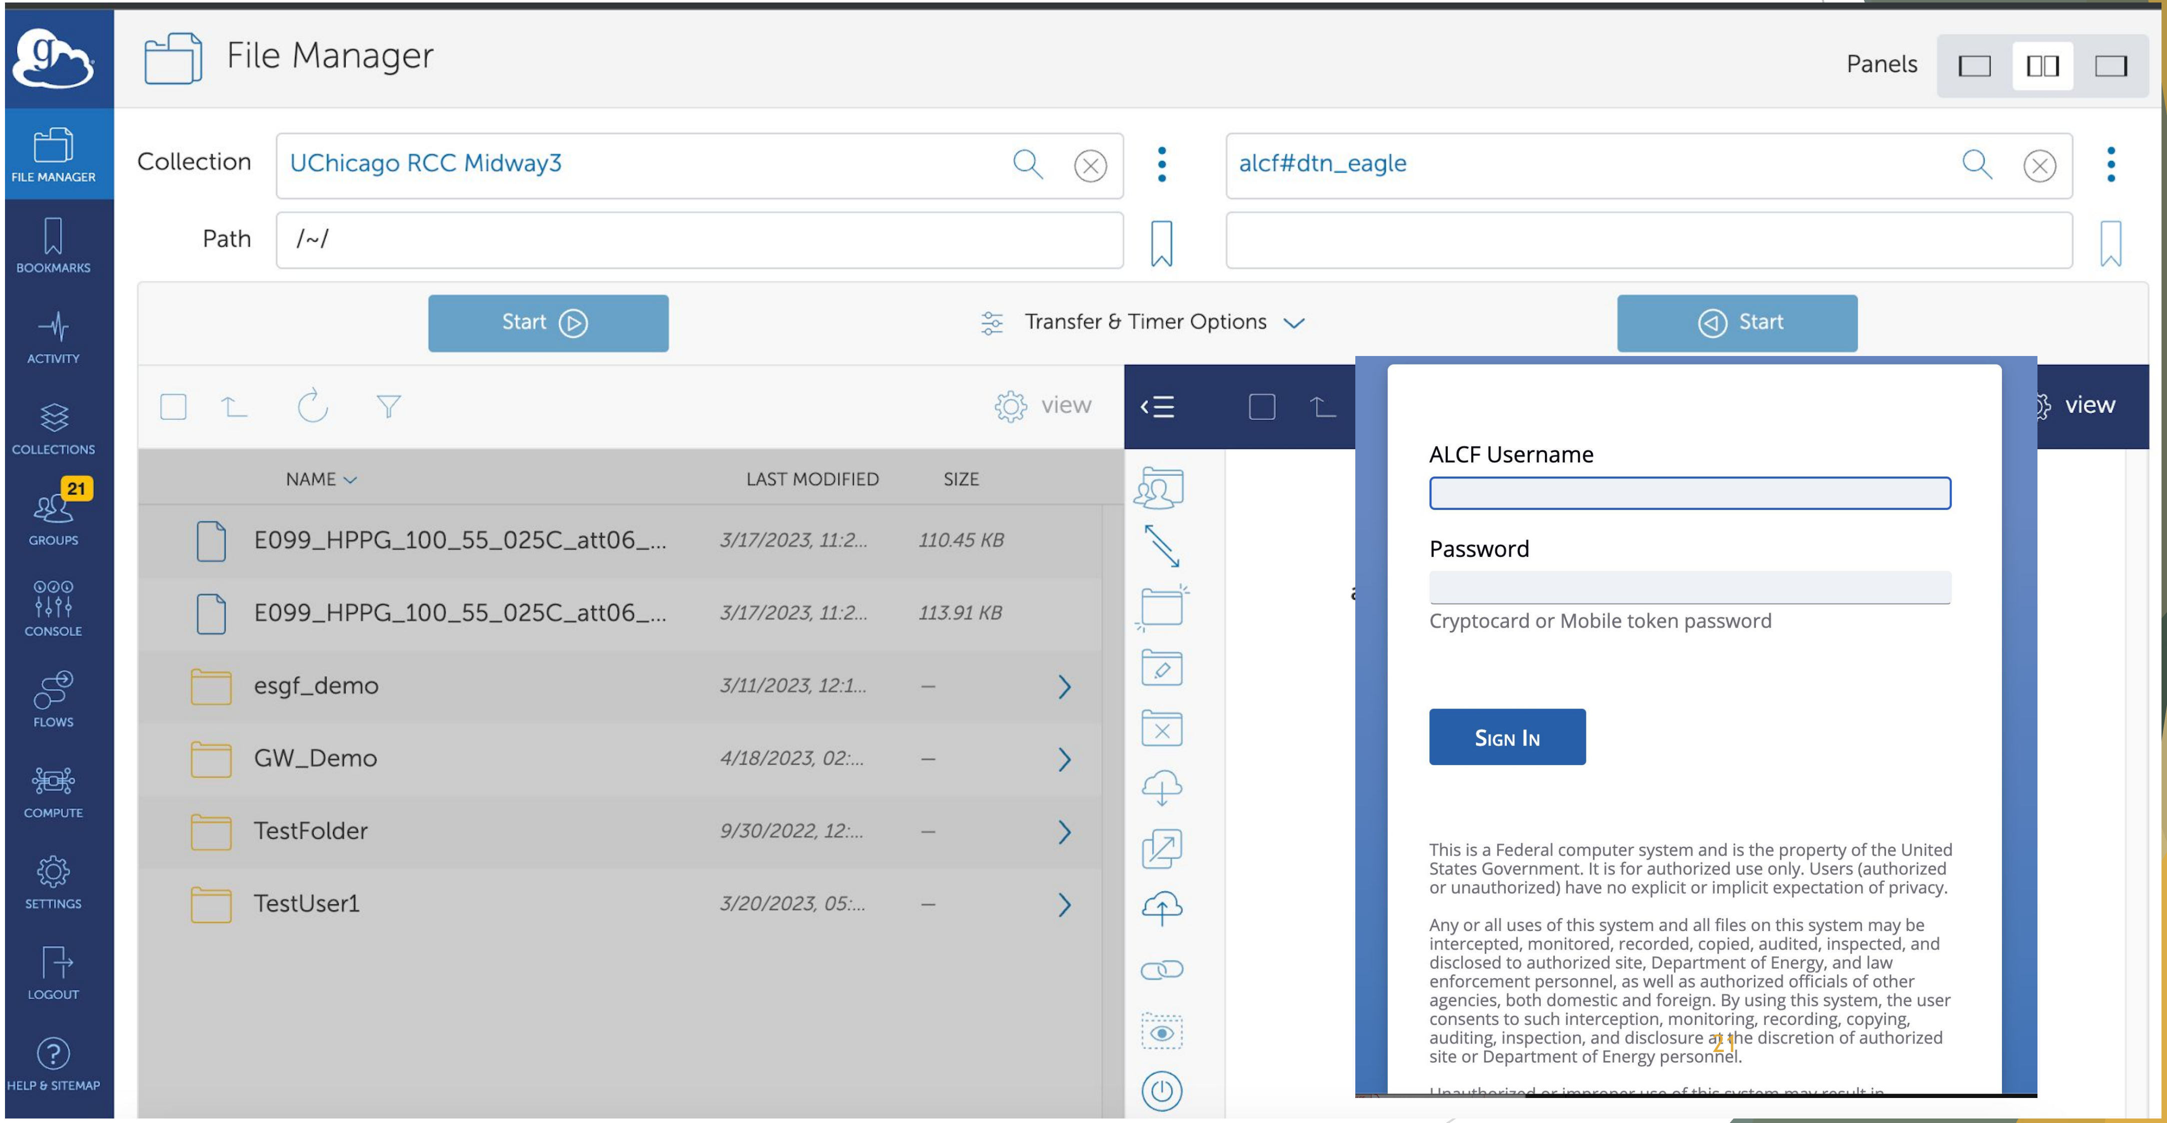

Collection: UChicago RCC Midway3

Path: /~/

Collection: alcf#dtn_eagle

| NAME | LAST MODIFIED | SIZE |
|---|---|---|
| E099_HPPG_100_55_025C_att06_... | 3/17/2023, 11:2... | 110.45 KB |
| E099_HPPG_100_55_025C_att06_... | 3/17/2023, 11:2... | 113.91 KB |
| esgf_demo | 3/11/2023, 12:1... | – |
| GW_Demo | 4/18/2023, 02:... | – |
| TestFolder | 9/30/2022, 12:... | – |
| TestUser1 | 3/20/2023, 05:... | – |

ALCF Username

Password

Cryptocard or Mobile token password

This is a Federal computer system and is the property of the United States Government. It is for authorized use only. Users (authorized or unauthorized) have no explicit or implicit expectation of privacy.

Any or all uses of this system and all files on this system may be intercepted, monitored, recorded, copied, audited, inspected, and disclosed to authorized site, Department of Energy, and law enforcement personnel, as well as authorized officials of other agencies, both domestic and foreign. By using this system, the user consents to such interception, monitoring, recording, copying, auditing, inspection, and disclosure at the discretion of authorized site or Department of Energy personnel.

Unauthorized or improper use of this system may result in

# Transfer/sync options

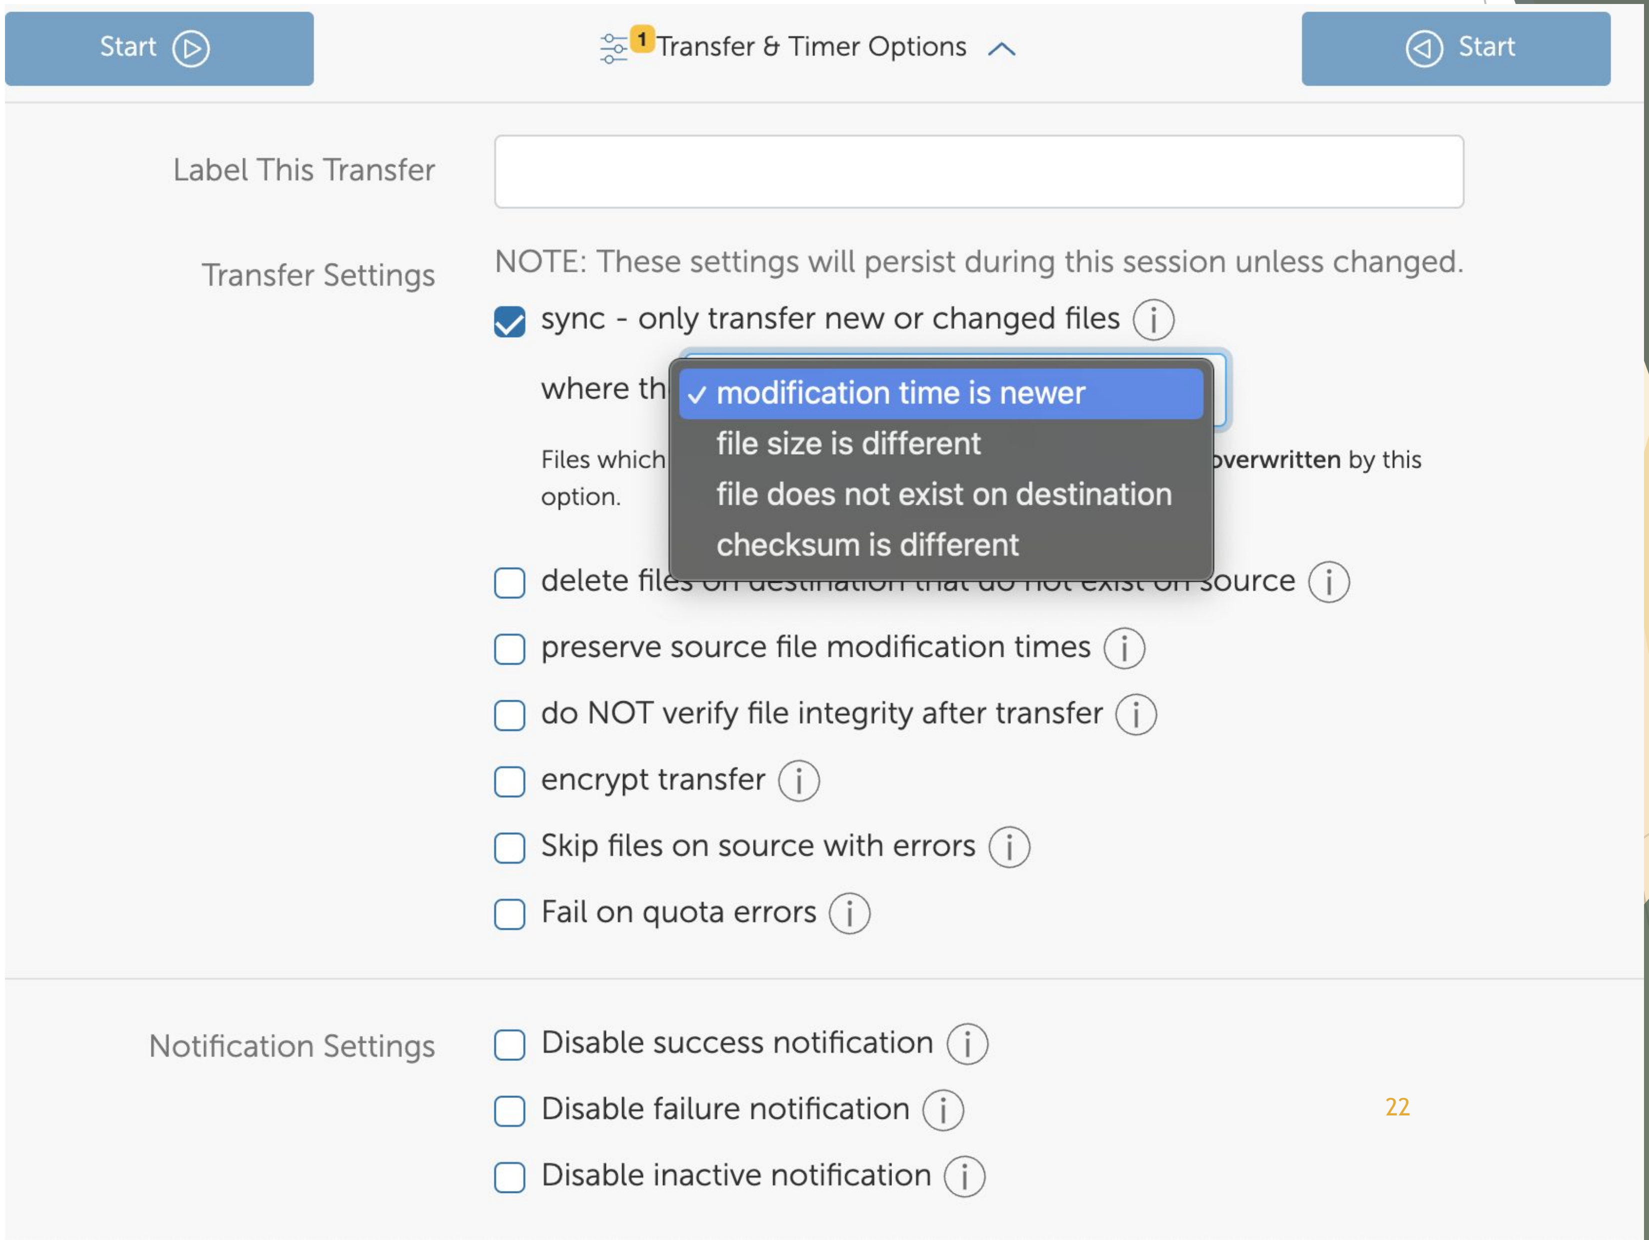

Label This Transfer

Transfer Settings

NOTE: These settings will persist during this session unless changed.

☑ sync - only transfer new or changed files (i)

where th ✓ modification time is newer

- file size is different
- file does not exist on destination
- checksum is different

Files which overwritten by this option.

□ delete files on destination that do not exist on source (i)

□ preserve source file modification times (i)

□ do NOT verify file integrity after transfer (i)

□ encrypt transfer (i)

□ Skip files on source with errors (i)

□ Fail on quota errors (i)

Notification Settings

□ Disable success notification (i)

□ Disable failure notification (i)

□ Disable inactive notification (i)

# Globus provides Fast, reliable file transfer …from any to any system

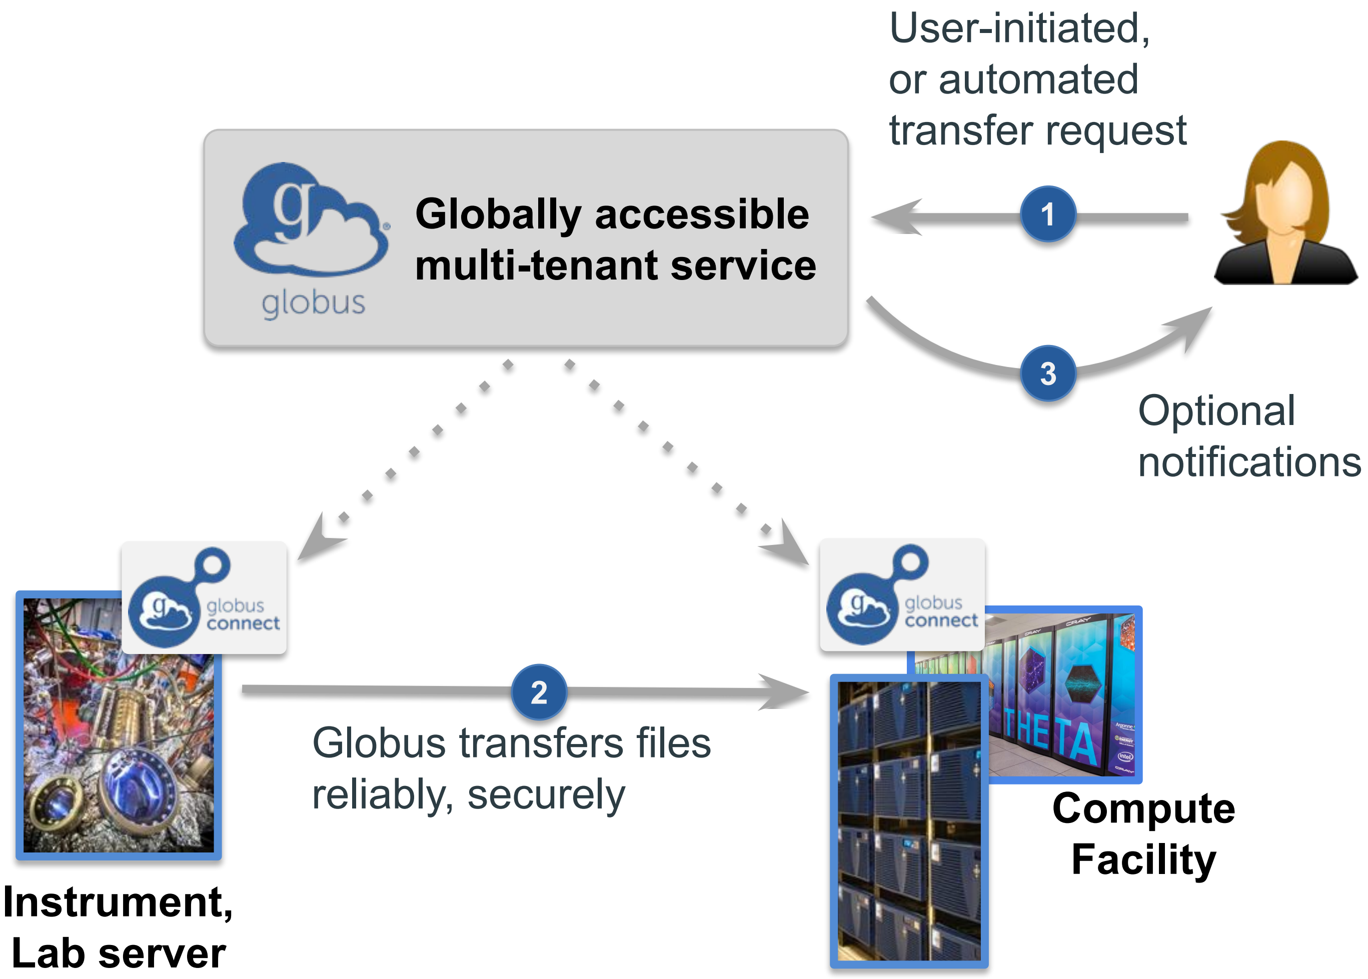

- Fire-and-forget transfers/sync
- Optimized speed
- Assured reliability
- Unified view of storage
- HTTP/S access to data

# Globus goes on the road

- Upload photos from mobile device
- Leverages HTTP/S upload and responsive web application

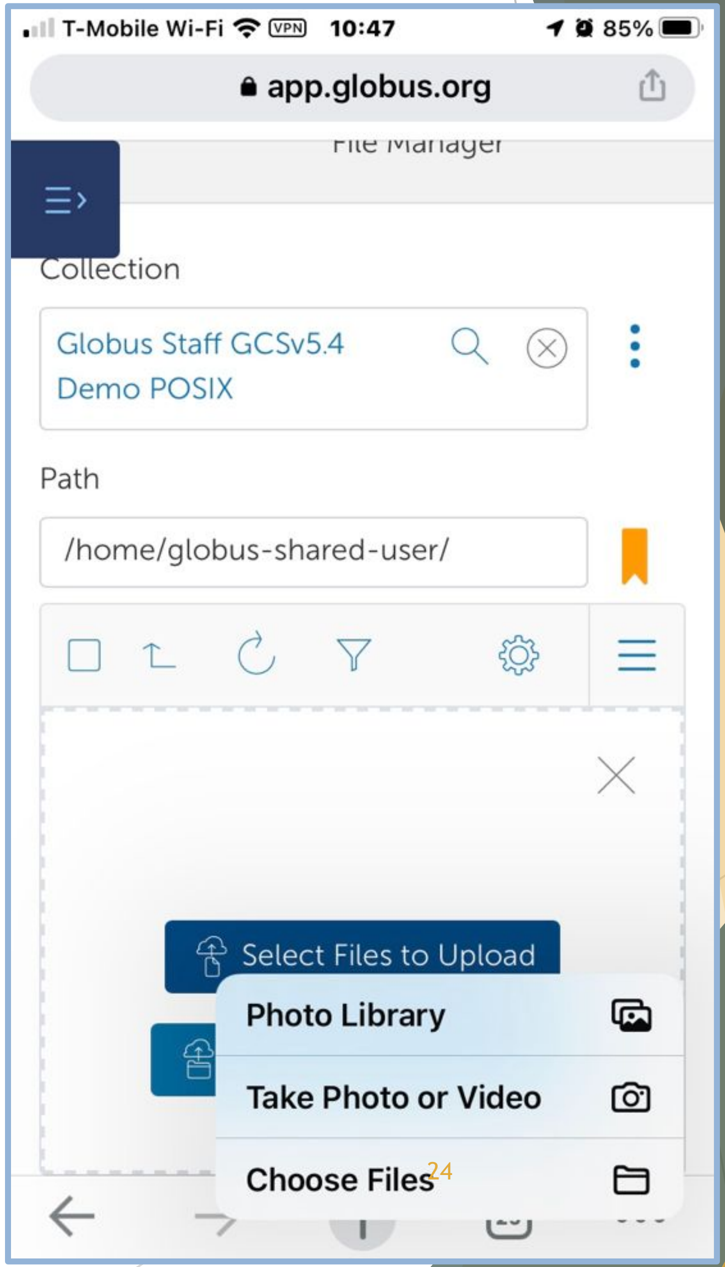

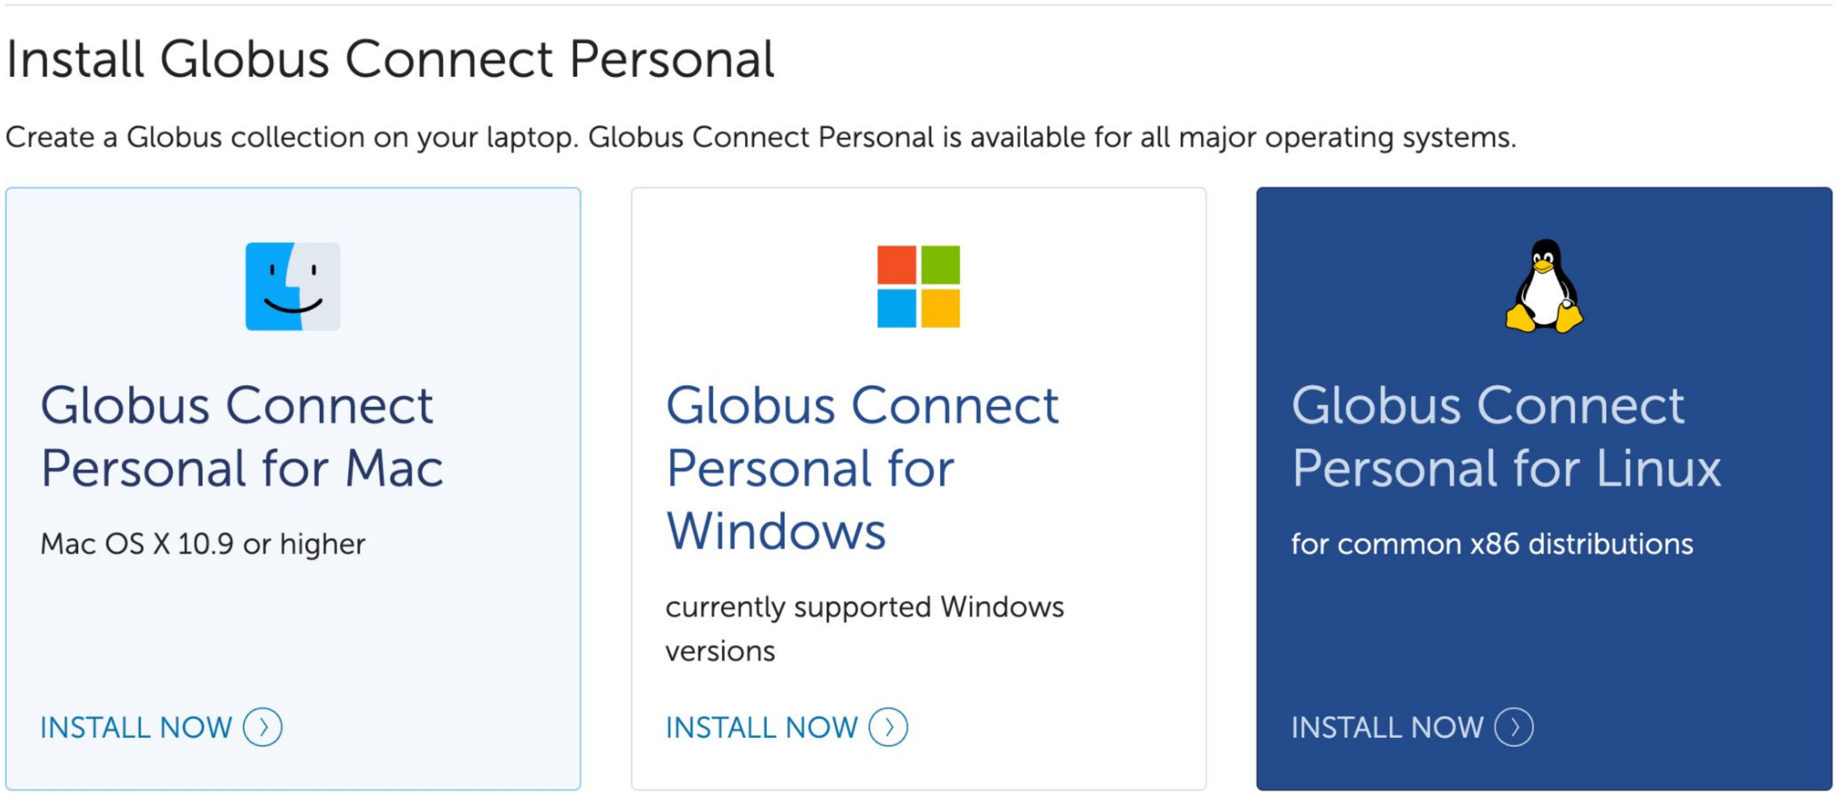

Install Globus Connect Personal

Create a Globus collection on your laptop. Globus Connect Personal is available for all major operating systems.

Globus Connect Personal for Mac

Mac OS X 10.9 or higher

INSTALL NOW ›

Globus Connect Personal for Windows

currently supported Windows versions

INSTALL NOW ›

Globus Connect Personal for Linux

for common x86 distributions

INSTALL NOW ›

## Globus Connect Personal

- ➢ Free Clients to easily and reliably move and share data from your personal computer or laptop to interact with other Globus collections.
- ➢ Easily download data from the cloud or campus computing cluster on to your laptop

# What current resources are available to researchers on the PA Science DMZ?

# Compute/Storage Options via the PA Science DMZ

**Over Internet2**

- Anvil (RCAC, Purdue)
- Delta (NCSA)
- Expanse (SDSC)
- UChicago AI Cluster
- Midway (RCC, UChicago)
- Kubernetes Clusters
- Polaris (ALCF)
- Perlmutter (NERSC)
- Frontera (TACC)
- Bebop (LCRC, ANL)
- Bridges-2 (PSC)
- FASTER (TAMU)

**Over Internet2:**

- Internet2 Cloud Connect -
  - AWS, Google,Azure, Oracle
- Open Science Grid
- National Research Platform

**Over DOE’s ESnet:**

- ALCF-Polaris
- NERSC- Perlmutter
- Bebop (LCRC, ANL)
- Frontier, Summit, Quantum- OLCF

Pennsylvania
healthcare
Science DMZ
serve
programs
University
events
kinber
I2PX
cost-effective
relationship
provides
service
community
government
Intermediate Unit
libraries
high-speed
vision
cybersecurity
support
Workforce Development
mission
unparalleled
strategic
leverages
network
Internet2
managed services
corporations
empowering
interest
digital technologies
hub
Community College
press
driving
KeystoneREN
DTN
HBC
progress
foster
research
Open Science Grid
government
state
cloud
cyberinfrastructure
policy
access
connect
education
Globus
secure
Artificial Intelligence
promotes
under resourced
solutions
for-profit
focus
public
areas
partnerships
usability
Career Technology Center
Broader Impacts
acceptable
safety
networks
centers
affiliate
affiliations
advanced
educational
operations
K12
digital equity
networking
excited
perfSONAR
collaboration
advance
innovative
connectivity
entities
institutions
non-profit
ESnet
services
underserved

"

# NETWORK AS ~~INFRASTRUCTURE~~ AN INSTRUMENT

"

KeystoneREN is the data circulatory system of Research and Education within Pennsylvania connecting users to resources, collaborators, and the world.

Ken Miller – ken@keystoneren.org

# Appendix

# Who is KeystoneREN?

Keystone REN, LLC, Lititz, PA, is a subsidiary of KINBER. KINBER, a non-profit, works with communities, governments, businesses, and other non-profits to drive solutions that support digital equity and inclusion.

The driving focus of KeystoneREN is to advance research and education networks and bring connectivity to underserved areas, empowering communities across the state. Our core competency is advanced networking and R&E cyberinfrastructure.

# KINBER/PennREN → KeystoneREN

- KINBER was founded in 2010.
- Firstlight largely bought all of KINBER's PennREN fiber and networking assets along with commodity customers on May 1, 2021. KINBER retained existing the Internet2 customers as well as the Internet2 Connector status.
- KINBER established Keystone REN LLC, as a non-profit LLC, on August 8, 2023
  - KeystoneREN remains the only statewide research and education network across the state of Pennsylvania
- Grant Dull, previously of KINBER and FirstLight, was hired as the Keystone Executive Director on July 24, 2023.
- Ken Miller, previously of ESnet and Penn State, was hires at the Chief Technology Officer on August 26, 2024.
- Jennifer Oxenford, previously of NYSERnet and KINBER, was hired as the Chief Relationship Officer on October 1, 2024.

# KINBER/PennREN services → KeystoneREN services

### KeystoneREN

Internet2 R&E ←

Internet2 Peering I2PX ←

Internet2 Cloud Connect ←

Keystone Member Exchange ←

### KINBER

~~Internet2 R&E~~

~~Internet2 Peering IX~~

~~Internet2 Cloud Connect~~

~~Keystone Member Exchange~~

~~PennREN Fiber and Network Assets~~

~~PennREN IP Space and ASN~~

~~PennREN Commodity Customers~~

~~PennREN Managed Routers~~

Digital Inclusion

Digital Equity

### FirstLight

→PennREN Fiber and Network Assets

→ PennREN IP Space and ASN

→PennREN Commodity Customers

→PennREN Managed Routers

# KeystoneREN advantages over previous network

- No longer locked into Crown Castle maintenance agreement on fiber operations and maintenance
- No longer burdened with high fiber plant operating costs reducing overhead
- Leveraging wholesale circuit procurement drive down customer costs
- Next generation 400G capable equipment provide greater operational efficiencies
- This sustainable network can be scaled as the growth scales.

# KeystoneREN Network Diagram

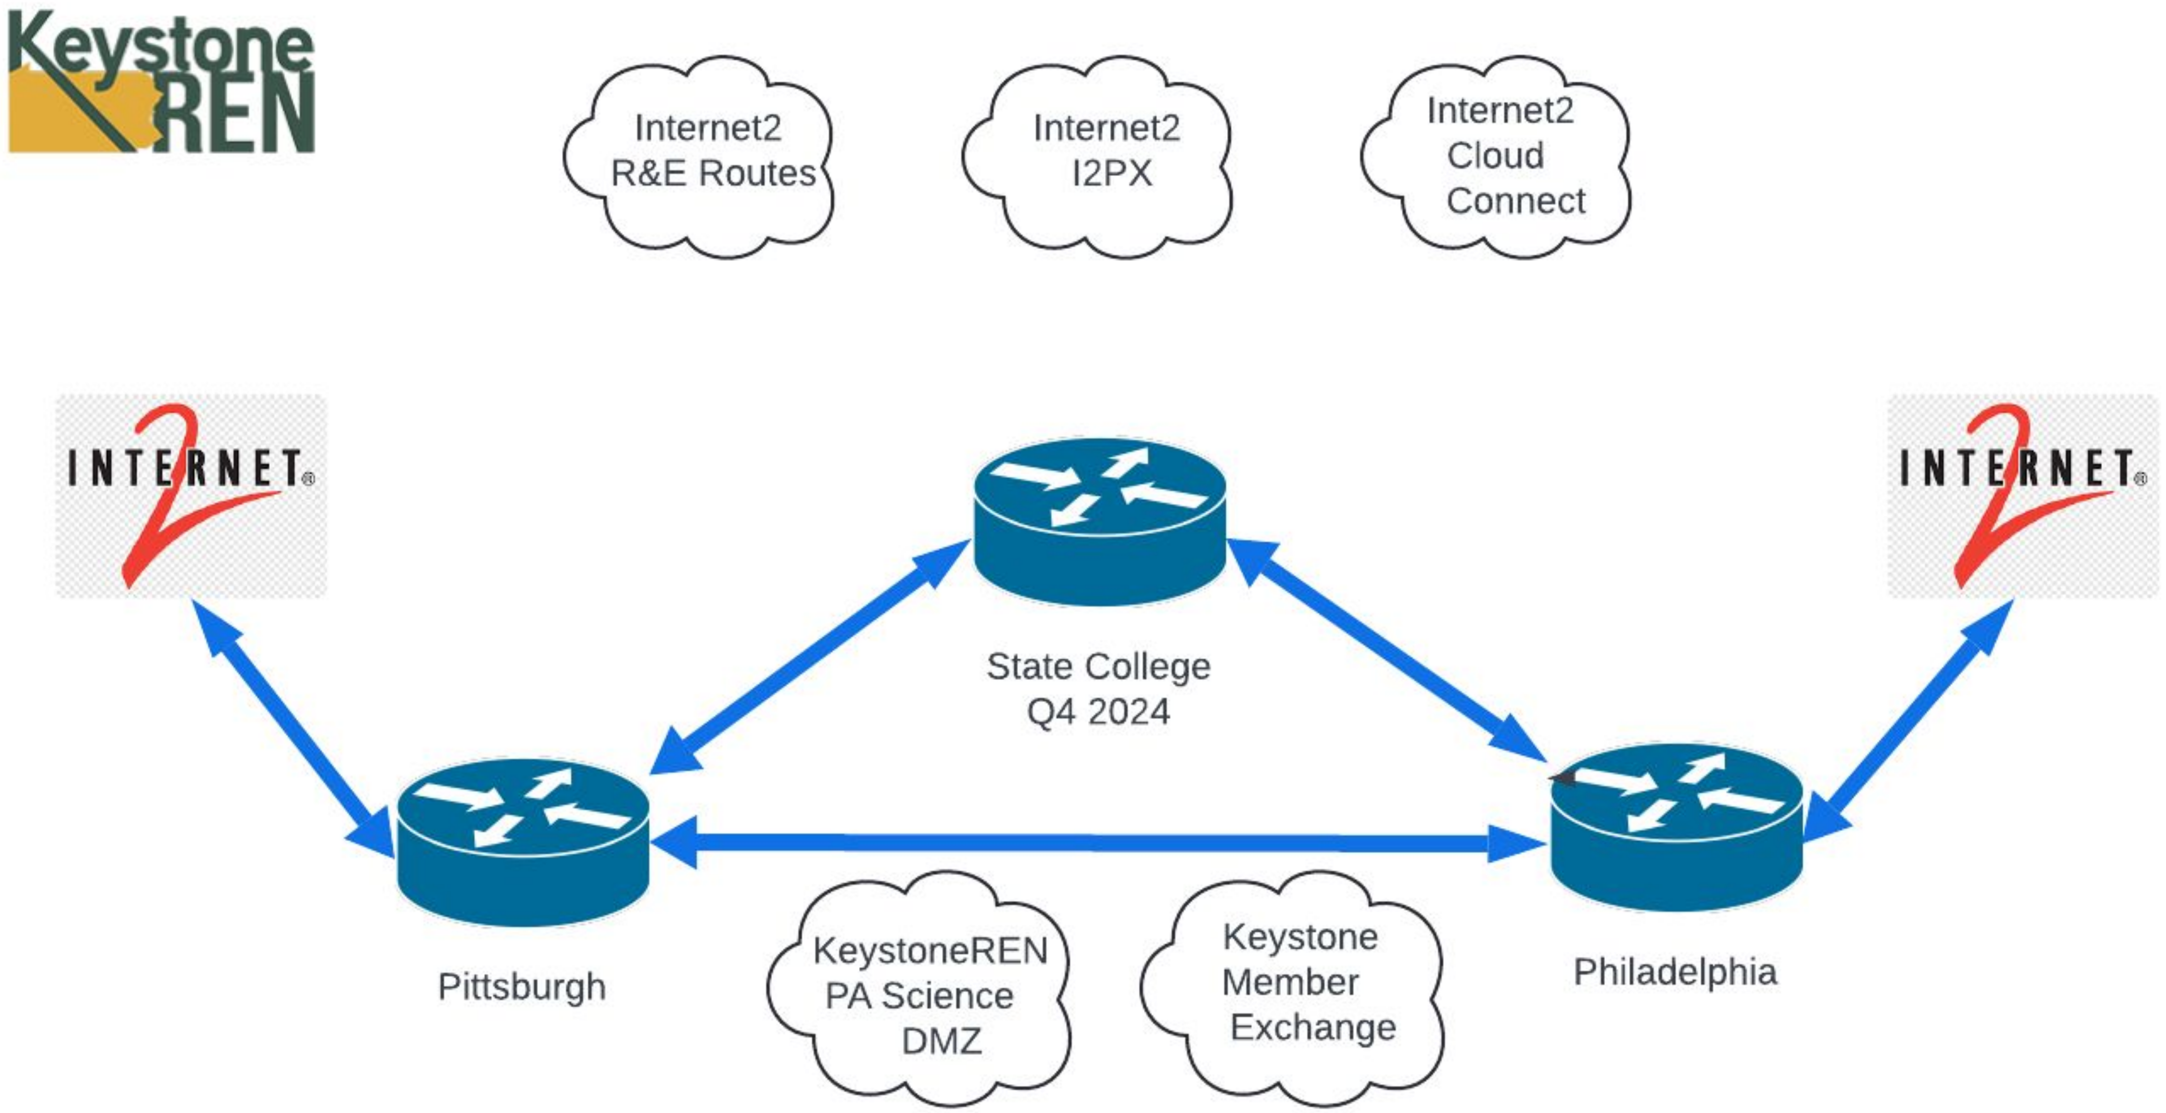

# NETWORK DIAGRAM

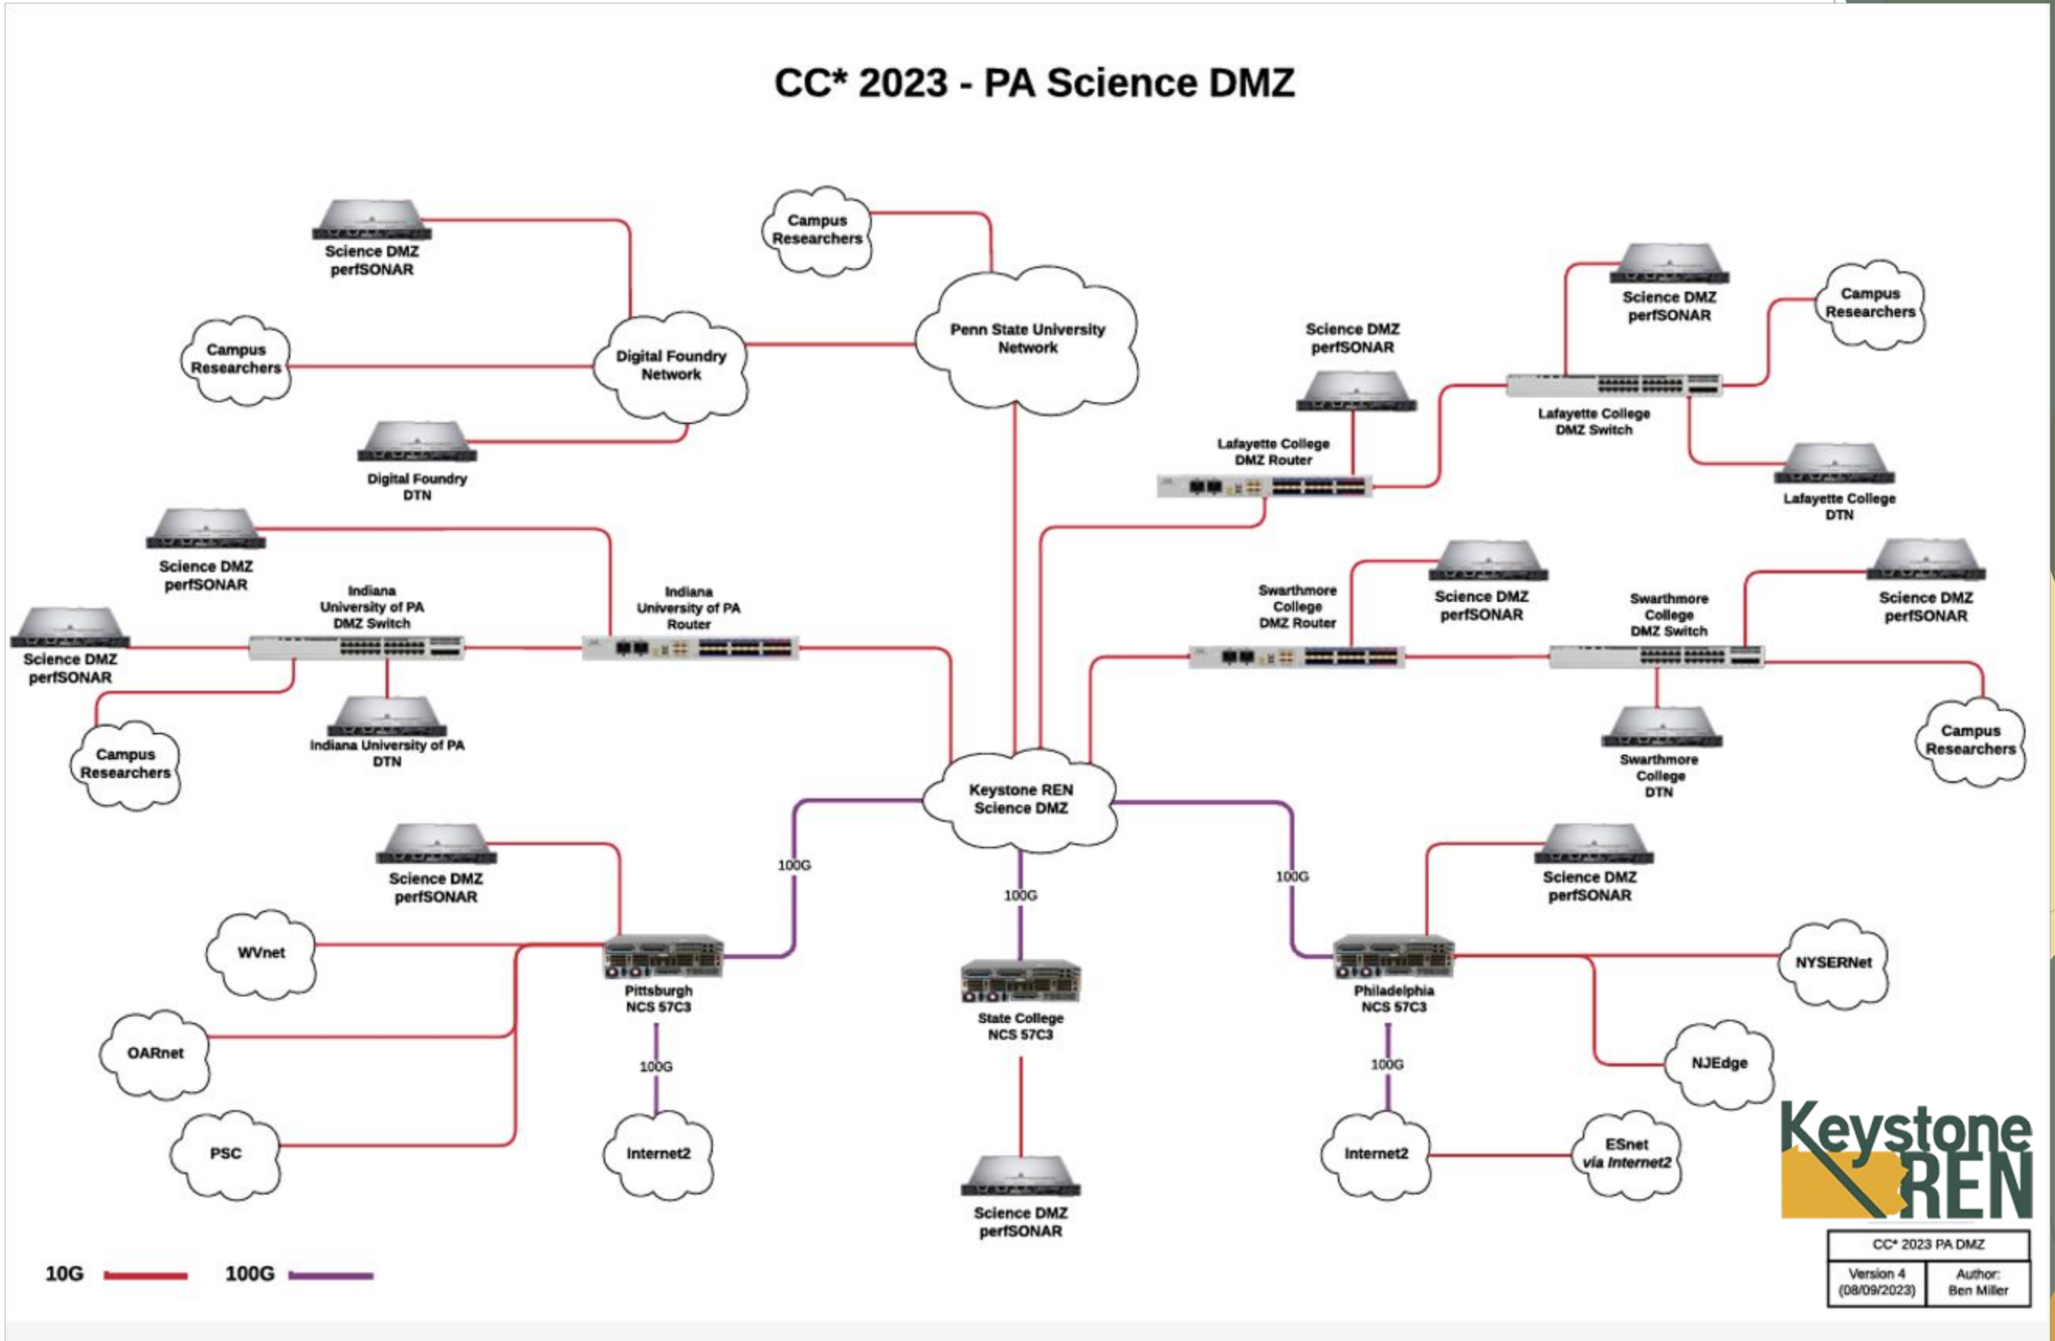

CC* 2023 - PA Science DMZ

# CAMPUS NETWORK DIAGRAM

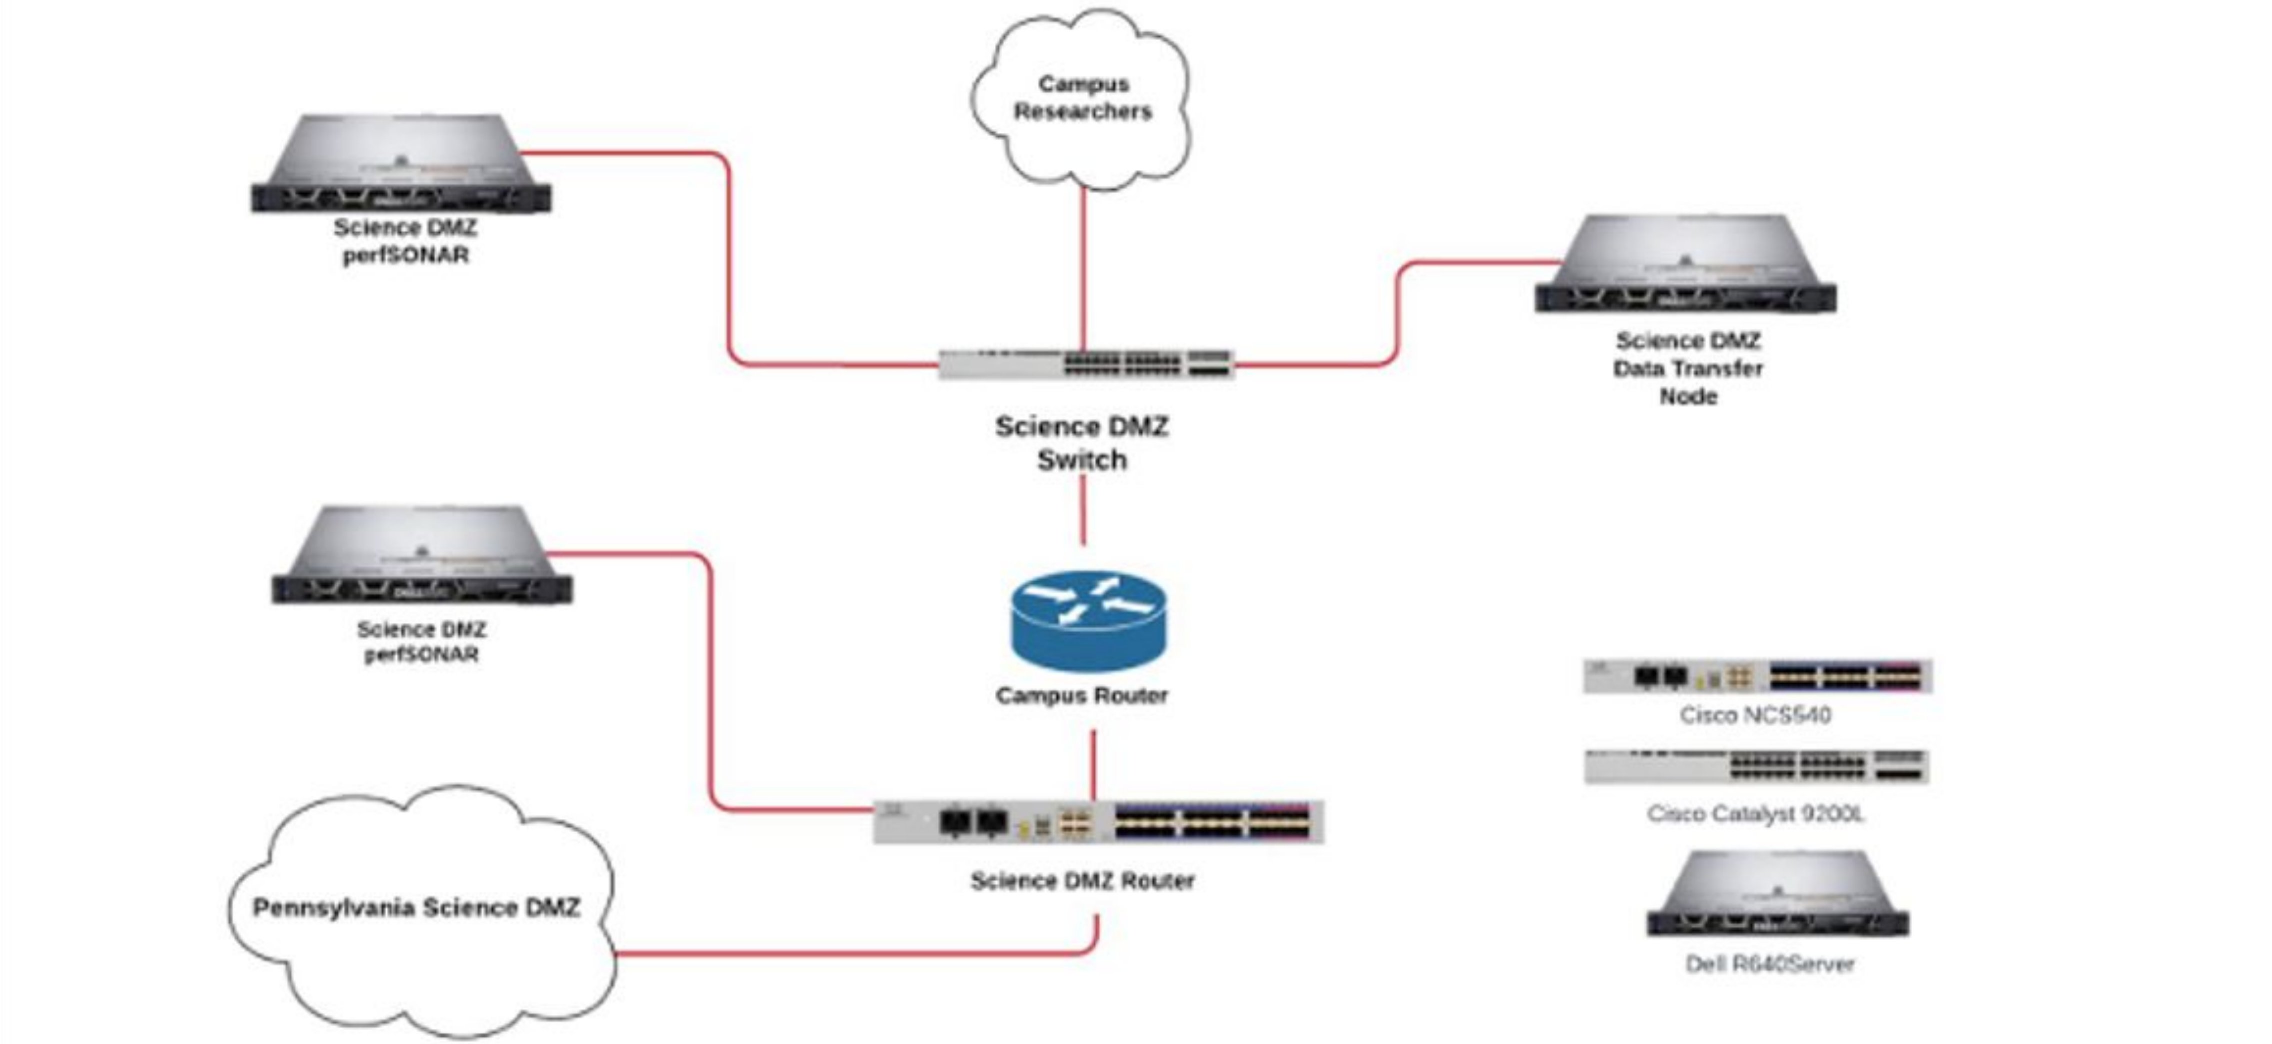

*Figure 3 - Campus Science DMZ*

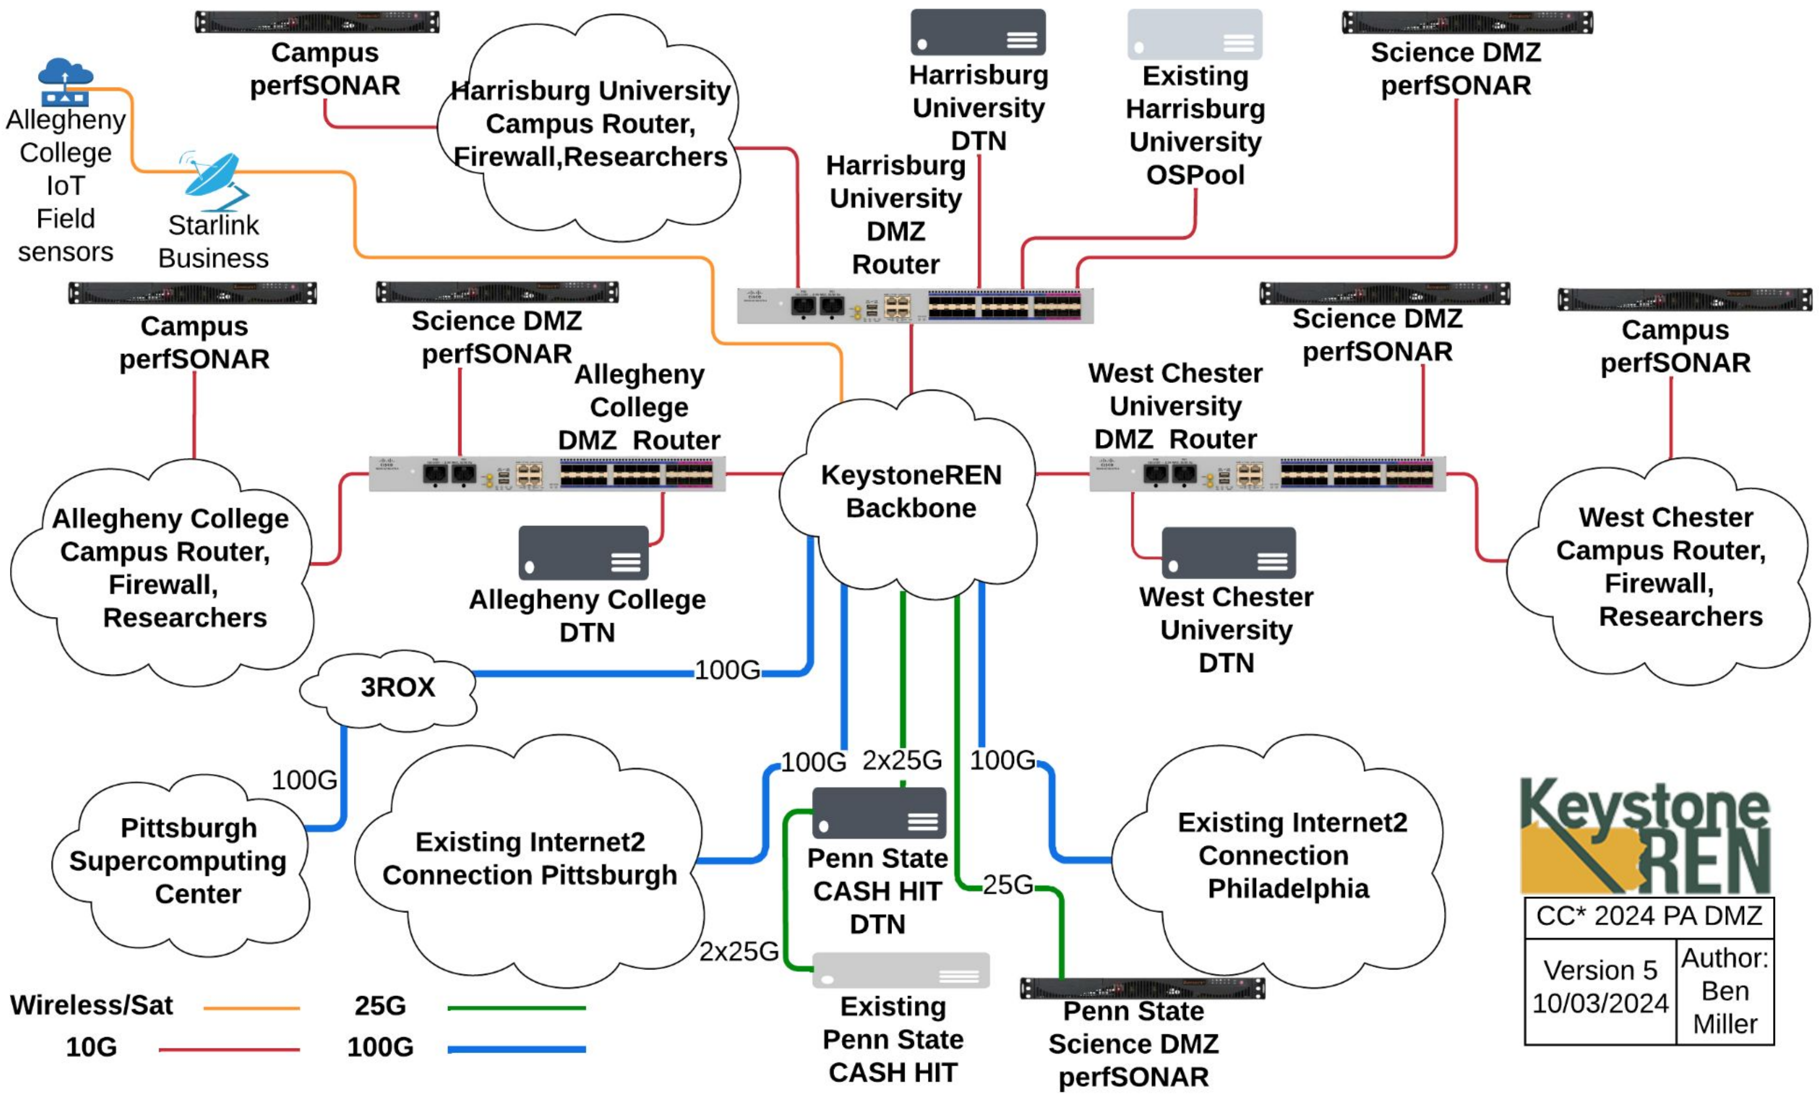

Allegheny College IoT Field sensors

Starlink Business

Campus perfSONAR

Harrisburg University Campus Router, Firewall,Researchers

Harrisburg University DTN

Existing Harrisburg University OSPool

Science DMZ perfSONAR

Harrisburg University DMZ Router

Campus perfSONAR

Science DMZ perfSONAR

Allegheny College DMZ Router

KeystoneREN Backbone

West Chester University DMZ Router

Science DMZ perfSONAR

Campus perfSONAR

Allegheny College Campus Router, Firewall, Researchers

Allegheny College DTN

West Chester University DTN

West Chester Campus Router, Firewall, Researchers

3ROX

100G

100G

Pittsburgh Supercomputing Center

Existing Internet2 Connection Pittsburgh

100G

2x25G

100G

Penn State CASH HIT DTN

25G

Existing Internet2 Connection Philadelphia

2x25G

Existing Penn State CASH HIT

Penn State Science DMZ perfSONAR

Keystone REN

CC* 2024 PA DMZ

Version 5
10/03/2024

Author: Ben Miller

| | |
|---|---|
| Wireless/Sat ━ | 25G ━ |
| 10G ━ | 100G ━ |

**Figure 2** – PA-Science-DMZ-2024 Technical Diagram