

Engagement and Performance Operations Center (EPOC)

(Formerly known as ReSEC)

NSF Award #1826994

Year 4 Quarter 4 and Annual

1 April 2021 through 31 Mar 2022

PIs: Jennifer Schopf (IU), Jason Zurawski (LBNL), Dave Jent (IU)

Summary

The goal of the EPOC project is to provide researchers and network engineers with a holistic set of tools and services needed to debug performance issues and enable reliable and robust data transfers. It supports six main activities: Roadside Assistance and Consulting, Application Deep Dives, Network Analysis using NetSage, Data Mobility Exhibition, the provision of Managed Services, and Training. Year 4 highlights include:

- Two new Regional Partners: NOAA N-Wave and University of Hawaii's Pacific Islands Research and Education Network (PIREN) as Regional Partners,
- Sixty-three completed and twenty-five ongoing Consultations,
- Four completed Deep Dive reports and seven more completed but waiting on institutional final approval for reports to be published, and
- Significant joint training.

1. EPOC Overview

The Engagement and Performance Operations Center (EPOC) is a production platform for operations, applied training, monitoring, and research and education support. EPOC is a collaborative focal point for operational expertise and analysis and is jointly led by Indiana University (IU) and the Energy Sciences Network (ESnet). EPOC provides researchers and network engineers with a holistic set of tools and services needed to debug performance issues and enable reliable and robust data transfers. By considering the full end-to-end data movement pipeline, EPOC is uniquely able to support collaborative science, allowing researchers to make the most effective use of shared data, computing, and storage resources to accelerate the discovery process.

EPOC supports six main activities:

- Roadside Assistance and Consulting via a coordinated Operations Center to resolve network performance problems with end-to-end data transfers reactively.
- Application Deep Dives to work more closely with application communities to proactively understand full workflows for diverse research teams in order to evaluate bottlenecks and potential capacity issues.

- Network Analysis enabled by the NetSage monitoring suite to proactively discover and resolve performance issues.
- The Data Mobility Exhibition/Baseline Performance Testing to test transfer times against known “good” test points, with the goal of transferring a TeraByte in an Hour.
- Provision of Managed Services via support through the Regional Partners.
- Training to ensure effective use of network tools and science support.

This report details the staffing, collaboration, and focused work in each of the six activities during Year 4. Meetings and publications are referred to throughout using [#], where the number is the item listed in Section 12 and 13. Note that at the time of this report, COVID-19 and its associated travel restrictions were in a state of high fluctuation, so many meetings were virtual.

2. Staffing and Internal Coordination

For Year 4, funded project staff included:

- Jennifer Schopf, IU, PI - Overall Project Director
- Jason Zurawski, LBNL, co-PI, Deep Dives and Managed Services Lead
- Hans Addleman, IU, Roadside Assistance and Consulting Lead
- Dan Doyle, IU, System Architect - Measurement and Monitoring co-Lead
- Heather Hubbard, IU, Project Coordination
- Brenna Meade, IU, Roadside Assistance
- Ken Miller, LBNL, Science Engagement and DME Lead
- Ed Moynihan, IU, Science Engagement
- George Robb, LBNL, Managed Services Support
- Doug Southworth, IU, Partner Coordination and Deep Dive Support

The only Year 4 addition was Brenna Meade, IU who started in June 2021 and focuses on Roadside Assistance and some of the broader outcomes for that activity.

Dave Jent is a co-PI, but due to his position at IU is unable to formally charge the project. The IU Global NOC Software team provides 0.5 FTE of developer support for the NetSage deployments.

3. Project Partners

EPOC has three types of partners: *Regional Partners*, who are deploying the infrastructure EPOC supports and use their members for outreach for EPOC, *Infrastructure Partners*, who are themselves collaborations that support a variety of cyberinfrastructure (CI) services, and *Science Community Partners*.

3.A Regional Partners

During Quarter 4, EPOC welcomed two new partners. The full list includes:

- **Front Range GigaPop (FRGP)**, the regional collaboration of networks that cover the western states, including Colorado, Wyoming, Arizona, Idaho, Utah, and New Mexico.
- **The Great Plains Network (GPN)**, the regional network that serves North Dakota, South Dakota, Nebraska, Iowa, Minnesota, Kansas, Missouri, Oklahoma, and Arkansas.
- **iLight**, the regional network for Indiana.
- **The Keystone Initiative for Network Based Education and Research (KINBER)**, the regional network for Pennsylvania.
- **The Lonestar Education and Research Network (LEARN)**, the regional network for Texas.
- **National Oceanic and Atmospheric Administration (NOAA) N-Wave**, the R&E network for NOAA. *Joined Year 4 Quarter 1.*
- **The Ohio Academic Resources Network (OARnet)**, the regional network for Ohio.
- **Pacific Islands Research and Education Network (PIREN) and the University of Hawai'i System Network**, which provides R&E network capacity to interconnect Pacific Islands with each other and the global R&E network fabric, including links to Australia and Guam, in addition to connectivity for the University of Hawai'i system and Mauna Kea and Haleakala astronomy observatories. *Joined Year 4 Quarter 3.*
- **Pacific NorthWest GigaPop (PNWGP)**, which provides access to next generation internet services and technologies throughout the Pacific Rim, but in the US primarily in California, Oregon, and Washington State.
- **Southern Crossroads (SoX)**, the regional network for much of the southeastern part of the US, including parts of Alabama, Georgia, South Carolina, and Tennessee.
- **Sun Corridor Network (SCN)**, the regional network for the state of Arizona.
- **Texas Advanced Computing Center (TACC)** at the University of Texas at Austin, United States, an advanced computing research center

A complete listing of our activities with these partners is given in Section 13.

In addition, as we formalize our Regional Partners program, discussions with NYSERNet, the R&E network for the state of New York, are ongoing.

In August, we held strategic planning meetings with eight of the Regional Partners, which resulted in an extensive list of action items, also detailed in Section 13. At a higher level, we refreshed contact points, evaluated the use of the Roadside Assistance and NetSage with the teams, promoted the use of known good perfSONAR endpoints (and began documenting them),

and agreed to give several presentations in the upcoming months for partner members. We also encouraged them to peer with Route Views and offered assistance to do so.

3.B Infrastructure Partners

EPOC's Infrastructure Partners are used to leverage the different kinds of support offered by each group to expand the set of services available to the broader community. The current set of Infrastructure Partner organizations includes:

- **The Campus Research Computing Consortium (CaRCC)** is a consortium of over 30 campuses that facilitates access to cyberinfrastructure. Schopf is the main contact for this group. She is currently following many of their mailing lists and regularly attends several of the Track meetings, including the Emerging Centers Track. She attended PEARC'21 [20] as well as the Internet2/CaRCC Online Supporting Data-Driven Discovery on Campus [27]. Eli Dart and Brenna Meade attended a CARCC Systems Facing meeting [40] and gave two talks describing different aspects of EPOC work for them [104, 105]. In Year 5, we plan to re-meet with the CARCC leadership to define clearer goals for additional interactions.
- **Trusted CI: The NSF Cybersecurity Center of Excellence** supports cybersecurity for NSF funded projects. Addleman is the main contact for Trusted CI. Addleman attended one meeting run by Trusted CI [37]. EPOC and Trusted CI staff are working together to produce best security practice documents and recommendations for Science DMZ architectures through a joint engagement with the University of Arkansas (Case #173). The final document, entitled "Science DMZ Secure High Performance Data Transfer" [118], was published in Quarter 4 and a blog post about the work was posted by Trusted CI [119]. We collaborated on Consultations #173 and 194.
- **Internet2** supports solving common technology challenges for their over 200 educational, research, and community members. Schopf is the main contact for this organization. In Quarter 1, she had a meeting with Dana Brunson to discuss how interactions could go forward, and presentations are being planned for the fall TechEx meeting. We are still working with Internet2 to understand interactions in lieu of the regular in person meetings. Throughout the year, we attended several Internet2 meetings and gave presentations including:
 - Internet2 Online session on Identifying Cyberinfrastructure Gaps [9]
 - Addleman attended the Internet2 Community Voices session on Network Security [24].
 - Southworth attended the Internet2 hosted perfSONAR meeting [43].
 - Schopf attended the Internet2/CaRCC Online Supporting Data-Driven Discovery on Campus [27].
 - Schopf, Addleman, Southworth, Meade, and Miller attended the Internet2 TechExtra Conference [45] and gave 4 invited talks [112, 113, 114, 115].

We also collaborated on consultations #175 and 212. In Year 5, we plan to continue to attend Internet2 meetings, and have submitted several talks to their upcoming conferences.

- **The Quilt** provides a central organization for networks to share the best practices to support end user science. Zurawski is the primary contact for the Quilt and has regular meetings with their organization. EPOC participates with the Quilt by:
 - Zurawski served on the advisor committee for the Quilt-sponsored virtual 2021 CC* PIs meeting which was attended by multiple EPOC staff [29], who spoke at multiple sessions for the meeting [95, 96, 97, 98].
 - EPOC staff attended the Quilt Fall Meeting [32].
 - EPOC staff attended the Quilt Spring Meeting [60]
 - EPOC had articles in both the 2021 Quilt Circle [78] and the 2022 Quilt Circle [117]
 - We are forming a strategic arc with Quilt technical representatives to participate in ongoing work on the Data Mobility Exhibition for the remainder of 2021 and into 2022.
 - As a follow-up from the Quilt annual meeting in September 2021, EPOC has updated online guidance for the Campus CI Plans, which are posted on <https://fasterdata.es.net>.
- **The Science Gateway Community Institute (SGCI)** provides best practice recommendations and support for scientists building and using data portals. EPOC is a formal partner of SGCI and recommends SGCI services when appropriate. There were no formal meetings or interactions with SGCI during Year 4. In Year 5, additional effort will be made to interact with SGCI and see how EPOC can help them meet their goals and deliverables.
- **The Extreme Science and Engineering Discovery Environment (XSEDE)** supports a single virtual system and CI expertise through the Campus Champions. Schopf is the primary contact for this group and is part of the XSEDE Advisory Board (XAB) and attended roughly monthly meetings for the XAB. She is involved in discussions with XSEDE leadership to understand how the XSEDE transition to ACCESS would fit with EPOC plans. She is also part of two successful proposals to the ACCESS program, including one which will deploy NetSage on all Access Resource Providers for additional data. The transition from XSEDE to ACCESS will take place in September 2022 and will be a focus for our joint work in Year 5.

We are working closely with the **Globus team** on the new test sites for the DME (Section 7) as well as new ESnet testing DTNs. EPOC staff gave a presentation on the Data Mobility Exhibition as part of their webinar series [25] and attended GlobusWorld 2021 [12], where we also presented on our joint work with Arecibo [77]. In Quarter 3 and 4, members of the EPOC team worked with Globus to containerize the Globus Connect Server via Docker (<https://github.com/globus/globus-connect-server-deploy>). In Quarter 4, ESnet deployed a DTN in Washington DC as part of the DME using this docker container model. We also collaborated with them on several Roadside Assistance and Consultation Cases (107, 147, 159, 223, 229, 231, and 234).

EPOC continues to discuss infrastructure partnership with the **FABRIC testbed**, as well as its international component, FAB. We will be supporting several of the experiments and making

sure the end-to-end paths for the testbed are running smoothly. In Year 4, EPOC staff attended 3 different FABRIC workshops [3, 30, 49] that discussed the network infrastructure, experiments, and FABRIC rack deployments. FABRIC plans to start bringing experiments between institutions online in 2023, and EPOC engineers will work to make sure performance is as expected

In Quarter 3, we began a collaboration with the CC* CIRA: **Shared Arkansas Research Plan for Community Cyber Infrastructure (SHARP CCI)** project (NSF #2126108), for whom we had previously supplied a Letter of Collaboration. Schopf attended their kick off meeting [42] and gave a presentation on EPOC and support for CC* CI Plans [107]. We have agreed to be on call for any needed interactions through Year 5.

EPOC continues to work with two regional collaboration teams with CC* funding:

- **SWEETER – South West Expertise in Expanding, Training, Education and Research** (NSF #1925764)
- **CC* CIRA: Building Research Innovation at Community Colleges (BRICCS)** (NSF #2019136).

Zurawski has attended their joint kickoff meeting [36] and has provided an EPOC overview [102]. Discussion on future engagement is ongoing.

3.C Science Community Partners

EPOC Science Community Partners each consist of a collaboration of scientists which we envisioned would enable us to scale our reach to larger community groups. The current Science Community Partners include:

- **The Earth Science Information Partners (ESIP)** is a consortium of over 180 members that provides a forum for the Earth science data and technology community. We monitored several mailing lists and are looking for collaboration opportunities. This quarter, Schopf met with several ESIP members at AGU [50], where she also left over 100 EPOC Informational cards at posters over the 4-day meeting. We continue to explore opportunities with the geoscience community in Year 5.
- **The University of Hawai'i System Astronomy Community** supports 15 facilities with hundreds of researchers and experiments every year. This is complementary to our support of PIREN as a Regional Partner. We are also coordinating on Consultation #121 and 208.
- **The Midwest Big Data Hub (MBDH)** supports the use of data for a variety of applications and end users across twelve states. Southworth was planning to attend the All Hubs Meeting [41], a conference which includes all of the regional hubs, in June 2020, but this meeting was canceled due to COVID related travel restrictions, and held virtually in October 2021. We have found it challenging to engage with this organization due to their continuing issues with leadership changes and will assess if continued engagement makes sense in Year 5.

3.D External Advisory Committee

After meeting with the majority of our partners in August 2021, we scheduled an External Advisory Meeting to take place in October to discuss ways of engaging more strategically with partners and the broader CI ecosystem, identifying venues to publish success stories, and soliciting ideas for follow-on funding. We gave an overview of the Year 3 successes (49 presentations or publications, 4 new Regional Partners, 97 Roadside Assistance and Consultation Cases with institutions in 28 states, including 21 EPSCOR jurisdictions, only 47 with previous partners, etc). We had a wide-ranging discussion over possible next steps EPOC could take as part of its evolution including:

- As a neutral third party, with NSF support, we can broker solutions to performance problems in a way that an individual institution cannot
- We are improving science across the US, and increasing national competitiveness
- An emphasis that the “Terabyte in an hour” metric could have additional national impact
- The need to share additional quotes with relevant NSF staff when institutions tell us “We didn’t know we could do this”
- Conflicting opinions on the CC* program, in part because as it matured many grants now seem to exaggerate without meeting their promise. EPOC could help with this aspect through use of the DME, for example.

We will be following up with the EAB on several of these aspects in Year 5.

3.E Large Scale Facility Work

While EPOC focuses primarily on work with small and medium sized institutions, we also work with several of the NSF Large Scale Facilities. We attended several NSF-sponsored meetings for large scale facilities [1, 56], the latter which was run by the new Cyberinfrastructure (CI) Center of Excellence for Navigating the Major Facilities Data Lifecycle (COMPASS). Additional work included:

- **Arecibo Observatory:** We continued our work with staff at the Arecibo Observatory, University of Central Florida (UCF), the Texas Advanced Computing Center (TACC), the University of Puerto Rico, Globus, and the CI CoE Pilot to transfer data from the Arecibo data store to TACC, tracked as Roadside Assistance Case #107. In Quarter 3, Arecibo staff reported that their primary spinning media, over 1 Petabyte, had been fully transferred to TACC. Additional effort continues to copy the full data set still on tape storage, with a predicted completion time of December 2022. A summary of the engagement has been posted on the EPOC website [106]. This engagement was a significant success for the community in that “Golden Copy” (only source) data for a 20-year astronomy instrument was made accessible to the broader community. Robb gave an invited keynote presentation on our work with Arecibo [77] at Globus World [12], and there were several press releases on the topic as well [71, 72, 73]. Robb attended a University of Central Florida IT Webinar [26] and presented about the Arecibo Observatory data movement work [90] to the audience who was in part responsible for the management of the center. And significantly, the multi-

organizational team involved in this work received an HPC Wire Readers Choice Award at SC'21 [44], with partners Arecibo Observatory, ESnet, Globus at the University of Chicago, Texas Advanced Computing Center (TACC), University of Central Florida, and COMPASS.

- **Event Horizon Telescope (EHT):** In October, 2020, a conversation started between various NSF IRNC network operators and the researchers who are part of the international Event Horizon Telescope (EHT) project. Throughout the year, EPOC engineers worked with staff at the Submillimeter Array (SMA) and James Clerk Maxwell Telescope (JCMT), and the two instruments in Hawaii, to baseline their network performance. EPOC staff also worked with Haystack to complete testing from the instrument in Spain. This is being tracked as Roadside Assistance Case #121.
- **Large Hadron Collider:** Moynihan, Addleman, and Southworth attended the LHCOPN-LHCONE meeting in October, where Southworth presented how NetSage could be used to solve their need for cross-site monitoring as part of the upcoming data mover challenges [11, 35, 63]. We are also working jointly with the University of Hawaii and ESnet to test data performance to an ESnet DTN at CERN for LHC traffic (Case #211).
- **National Radio Astronomy Observatory (NRAO):** We worked with the South African National Research Network (SANReN) to debug performance problems on downloads from the NRAO archives in Charlottesville, VA (Case #210).
- **Vera Rubin Observatory (VRO):** Moynihan is a member of the VRO-NET and is tracking their needs for data movement and ways that EPOC can help them as they ramp up for the telescope becoming operational. He attended their annual meeting [7] and offered EPOC help to ensure effective data transfers when the instrument was sharing data more broadly.

In Year 5, we will continue our support for these Large Scale Facilities and are planning to reach out to and engage with other facilities to learn more about their data transfer challenges and to see if EPOC can provide assistance.

3.F External Partners

In addition to the previously mentioned partners, the EPOC team is coordinating with additional groups:

- The “**Toward the National Research Platform**” (TNRP) project (NSF #1826967), led by Larry Smarr and Tom Defanti, is tasked by NSF to stay in coordination with EPOC as both teams support the other CC* awardees. We attended the PRP Capstone event as well [18].
- EPOC also participates in the **Global Research Platform (GRP)**, an international collaboration that aims to integrate international fabrics and distributed cyberinfrastructure to support data-intensive scientific workflows. We attended their annual meeting [31]
- We are working closely with members of the **University of South Carolina Cyber Training team** (NSF#1829698), as detailed in Section 9.

- EPOC is working with the **CI Engineering Community** (<https://www.es.net/science-engagement/ci-engineering-lunch-and-learn-series/>) to catalog the presented webinars from the Lunch and Learn series held from 2017 to present. To the end of the reporting period, 125 webinars have been held, and are available on the EPOC YouTube channel as of March 31, 2021: <https://www.youtube.com/c/EPOC-IU-ESnet>. This included several presentations by the EPOC team [100, 110]
- Our work with BGP is complemented by the **GNA-G / APAN Routing Working Group (RWG)** (<https://www.gna-g.net/join-working-group/gna-g-routing-wg/>), which Meade and Addleman co-lead, along with Warrick Mitchell, part of Australia's Academic and Research Network (AARNet). EPOC staff hosted BOFs about the RWG at TNC [17, 84] and APAN [22, 87]. We also attended several GNA-G meetings on behalf of the RWG [15, 62] and also presented to them [82, 128].
The RWG meets monthly and works on routing cases in between meetings. At the end of Year 4, they had examined 14 cases, and resolved 10 of them. The first official meeting of the RWG was held in June, 2021. EPOC staff have given presentations at RWG meetings, including when Schopf presented a NetSage overview [85]. In Quarter 4, the RWG began working on best BGP practice documents for R&E networking as a community including scalability, load balancing, BGP peering agreements within the R&E community, and how to best use BGP communities, which will be published as EPOC White Papers in Year 5.
- In Quarter 2, EPOC was invited by the **Pacific Rim Applications and Grid Middleware Assembly (PRAGMA)** to give a talk, as part of their Reconnect Webinar series, on the end-to-end performance work EPOC is doing in support of international collaborations. Moynihan gave the presentation and participated in a discussion on how EPOC could partner with PRAGMA on future activities [21, 84]. EPOC continues to participate in PRAGMA and will stay engaged in Year 5 as they further define the role they will play in the community.

3.G Support for Grant Submissions

Throughout the year, EPOC provided support for organizations submitting grant proposals to the NSF. In Quarter 1, a BoF was held as part of the Great Plains Network Annual Meeting [13] describing how EPOC can support submissions to the NSF CC* Solicitation [17]. In Quarter 2, we participated in the CC* PI meeting [29], which included EPOC staff leading sessions on DME [95], how to do Deep Dives [96], the Deep Dive and RA outcomes [97], and R&E Routing [98].

In total, EPOC provided support for 11 CC* proposals. These are tracked as Roadside Assistance Cases 105, 142, 177, 179, 181, 184, 192, 193, 200, 207, and 209. For these proposals, we wrote CC* Letters of Collaboration (LoC) for NJEdge, OneNet, the University of South Dakota (USD), and Kent State.

In Year 5, we will continue to provide support and guidance upon request for organizations needing assistance with grant proposals.

4. Roadside Assistance and Consulting

A key aspect of the EPOC project is the operations center and process pipeline for immediate help, referred to as Roadside Assistance and Consulting, which assists collaborators when data sharing failures occur. EPOC coordinates with the multiple domains and organizations involved to achieve a resolution. More information about the Roadside Assistance and Consulting process is available: <https://epoc.global/wp-content/uploads/Roadside-Assistance.pdf>. Hans Addleman is the lead for this effort.

4.A Roadside Assistance Cases

In Year 4, we had four closed Roadside Assistance Cases:

- **76 - National Center for Atmospheric Research (NCAR):** NCAR distributes scientific weather data from instruments and simulations to 250 consortium member institutions using the Local Data Manager (LDM). The LDM team was testing a new version of the file transfer protocol, based on UDP multicast, on a 5-site testbed that includes NCAR, University of Virginia (UVA), University of Wisconsin Madison (UW-M), University of Washington (UW) in Seattle, and University of California San Diego (UCSD). In the test setup, UCSD and NCAR had issues with packet loss when sending and receiving from the other three sites. UW-M, UVA, and UW are able to transmit and receive data as expected from each other. EPOC staff worked with engineers for each institution as well as a suite of regional network engineers associated with the end points. The issue was narrowed down to a single switch in the FRGP network in Denver, Colorado, which needed to be replaced. Delays in procurement and installation pushed this installation to June, 2021, and now multicast testing is working as expected with no packet loss. An EPOC White Paper for this work is now available [75].
- **107 - Arecibo:** University of Central Florida (UCF) requested EPOC staff to consult with the staff at the Arecibo Observatory about moving 2 petabytes of data off site to a Microsoft Azure Cloud storage instance and a subset of that data to UCF's Advanced Research Computing Center (ARCC) cluster for data processing. It was discovered during these discussions that the data at Arecibo was the only copy of over 20 years of observations, emphasizing the need to have a backup copy. On November 6, 2020, one of the main cables supporting the radio telescope failed, causing major damage to the dish, and endangering the data center at Arecibo. This work became prioritized and it was evaluated that making a copy to Azure would take almost 3 years. In late November, the Texas Advanced Computing Center (TACC) offered to store the data indefinitely. On December 1, 2020, a second cable failed causing a collapse and destruction of the telescope. The data center was not damaged during the collapse but the data movement was seen as even more critical. In December, 2020, a larger plan to move the data from Arecibo to TACC was created with input from TACC, University of Puerto Rico (UPR), Arecibo, UCF, Globus, the CI CCOE Pilot at ISI, and EPOC staff. Arecibo staff purchased several portable network attached storage appliances (NAS) and began a cycle of loading them at the Arecibo data center, transporting the appliances to the UPR campus, and then transferring the

data to TACC using Globus. Baseline performance testing with perfSONAR showed a clean path between the UPR campus and TACC. Data from NetSage and logs provided by TACC and Arecibo show that over 1PB of data has been moved from Arecibo to TACC. Data will continue to be moved and Arecibo staff have started moving data off of tape backup to spinning disk for transfer to TACC. TACC engineers continue to work with Arecibo staff to make the data available via a portal for access by researchers. This work resulted in an HPC Wire Readers Choice Award, presented at SC'21 [44]. An EPOC White Paper for this work is now available [106].

- **164 - Kyungpook National University (KNU):** During routine NetSage traffic analysis of the TransPAC-PacWave 100G circuit, EPOC engineers noticed a significant amount of traffic from Kyungpook University in South Korea to CERN in Geneva, Switzerland. Data transfer performance was slow, averaging 30Mbps over the course of several weeks. EPOC engineers reached out to engineers at the Korea Advanced Research Network (KOREN) to inform them of their findings, as well as offer assistance with improving transfer speeds or finding a more efficient route. KOREN engineers investigated the problem, and as a result moved the traffic from KOREN to Korea Research Environment Open Network (KREONET). This put the traffic on a more direct path which reduced the number of network hops and round trip time. As a result, the transfer performance increased. An EPOC White Paper for this work is now available [69].
- **165 - National Institute of Informatics (NII):** EPOC staff, utilizing NetSage for routine traffic analysis on the TransPAC-PacWave 100G circuit, noticed the erroneous routing of traffic from the National Institute of Informatics (NII), in Japan, to Pohang University, in South Korea, which was being detected on the TransPAC NetSage Sensor in Seattle. This traffic was being routed over the TransPAC-PacWave 100G circuit to the US and then on to Chicago, where it then picked up Internet2's peering with KREONET and returned back across the Pacific to South Korea. Traffic going in the other direction, however, was routed correctly across available links in the region. EPOC staff worked with KREONET engineers to adjust the BGP advertisements to the APAN-JP router in Hong Kong, resulting in traffic taking a much more efficient route within Asia. An EPOC White Paper for this work is now available [70].

In Year 4, we had one ongoing Roadside Assistance Case:

- **121 - Event Horizon Telescope (EHT):** EPOC staff are part of an ongoing conversation between the astronomy researchers working with the Event Horizon Telescope (EHT), and the researchers supporting the NSF-funded International Research Network Connections (IRNC). The EHT science use cases involve transporting terabytes of data from eleven telescopes in distributed locations to the MIT Haystack Observatory, in Westford, Massachusetts, for consolidation and analysis. The telescopes are capable of a maximum streaming rate of 64Gbps. In previous quarters, tests were completed between the NOEMA telescope in France and the 30mm telescope in Spain. It was discovered that a 1G bottleneck at the control computers in both Spain and France existed, and discussions began with the EHT group to upgrade these to remove the limitation. The EHT systems engineer and his team are currently evaluating upgrade

options and once the upgrades are complete re-engage EPOC to retest the paths to Haystack. Testing to the two Hawaiian instrument sites was ongoing throughout Year 4, however, progress has stalled. In Year 5, we hope to re-engage with both Hawaii and the European sites.

4.B Consultation Cases

In Year 4, we had 63 completed and 25 ongoing Consultations. The most common Consultation topics were data transfer performance issues, Science DMZ deployment, security and firewall issues, and routing problems.

Completed Consultations included:

- **63 - National Oceanic and Atmospheric Administration (NOAA):** NOAA staff reached out to EPOC to request help and materials to run their own Deep Dive. EPOC staff provided materials and guidance. EPOC staff presented virtually on the process during a NOAA meeting and the ticket was closed.
- **105 - Lafayette College (LC):** The manager of research and high performance computing at Lafayette College requested information about Science DMZs, DTN design, and data architecture best practices. The discussion encompassed data transfer testing, network performance testing, and Science DMZ design for a potential NSF CC* grant proposal.
- **113 - Texas A&M University (TAMU):** As part of release testing, NetSage staff found a seven-day transfer between Texas Advanced Computing Center (TACC) and TAMU that moved over 24 TB of data at a rate under 500Mbps. TAMU engineers reported that they made some software upgrades to their switches and firewalls in January 2021. NetSage shows that the transfer rate is variable between 500Mbps and 1.3Gbps after the changes. Early in Year 3 Quarter 4 during a routine check of NetSage, it was found that the transfer rates had again fallen under 500Mbps again. EPOC and TAMU Engineers discussed the performance drop at that time and TAMU engineers found no apparent changes in the network configuration. TAMU engineers installed perfSONAR nodes to help troubleshoot the issue and then considered this issue resolved.
- **117 - Allen Institute (AI):** EPOC staff consulted with staff at the Allen Institute regarding CI resources in the community and funding opportunities.
- **135 - Texas Advanced Computing Center (TACC):** TACC was in the process of a network refresh and the manager of Network Operations asked for EPOC assistance during the planning stage. The discussion centered around switching and firewall hardware architecture.
- **143 - Arizona State University (ASU):** ASU staff asked for assistance in designing the data architecture for an experimental facility for the Compact X-Ray Free Electron Laser (CXFEL). The instrument is similar to the Department of Energy Light Sources and will support researchers in fields such as material science. Staff members were evaluating options for an effective data architecture, including the network infrastructure, storage, computational power to calibrate and analyze samples, and the workflow tools to be

used to control the components. EPOC staff answered questions and provided information about best practices used by other facilities with similar instruments.

- **147 - South African Radio Astronomy Observatory (SARAO):** The SARAO, in collaboration with the South African National Research Network (SANReN), approached EPOC for assistance on a multi-continent performance problem when downloading data from the National Radio Astronomy Observatory (NRAO) in Charlottesville, VA. After several weeks of debugging, a number of problems were identified and corrected, including:
 - Network routing abnormalities were found and addressed in SANReN's network.
 - Firmware updates were performed on SANReN's routing hardware.
 - An MTU configuration problem was discovered and corrected on the campus network for the University of Virginia, the campus network connection to NRAO.EPOC staff also helped the group evaluate data transfer tools, such as Globus, that could speed up the overall data transfer.
- **152 - Brown University:** A senior network engineer at Brown University asked questions about adding a firewall to the path of an established Science DMZ. EPOC staff answered questions and shared best security practices. The Brown engineer reported they were able to change the design of their Science DMZ based on the discussion and remove the proposed firewall from the path.
- **153 - Saint Louis University (SLU):** EPOC staff met with the Director of Network Services at SLU to describe support and services that the EPOC project could provide.
- **154 - NYSERNet:** A network engineer at NYSERNet reached out to EPOC for an overview of NetSage. EPOC and NetSage staff demonstrated NetSage for NYSERNet staff and discussed a NetSage deployment in NYSERNet and the possibility of becoming an EPOC partner.
- **155 - Texas A&M University San Antonio (TAMUSA):** A researcher at TAMUSA contacted EPOC for assistance in designing a Science DMZ network. Initial discussions indicated that the proposed network was being designed without identifying specific scientific use cases. This further revealed several questions regarding specific access policies, usage patterns, and technical requirements that could not be answered without a deeper understanding of the research landscape. EPOC, working with LEARN, encouraged TAMUSA to adopt some tactics from the Deep Dive approach to understanding research use cases, to better address the concerns regarding how to design, implement, and operate a Science DMZ for the campus. TAMUSA is currently gathering this information and will use the outcomes to develop a strategic plan for design and implementation of the infrastructure.
- **157 - Syracuse University:** Syracuse University runs a DDoS scrubbing service at the border of their campus network that requires special configurations for packet size manipulation and redirection of traffic to the scrubbing services network. The Chief Network Architect at NYSERNet reached out to EPOC for help investigating the performance impact the device had on traffic on the campus network, as many of their connector universities were considering adding this type of service as well. EPOC engineers worked with the NYSERNet staff to test the effect this service had on data transfers using perfSONAR. The investigation found that modern DTN and perfSONAR

servers could handle the packet size manipulation with little to no impact at 10Gb/s speeds.

- **158 - New York University Langone Health (NYULH):** Staff at NYULH were investigating the Science DMZ model and how it might fit their network. A meeting was held to discuss best Science DMZ architecture and security practices.
- **159 - Yale:** The program director for the next generation network project at Yale reached out to discuss some performance abnormalities they experienced with Globus file transfers. EPOC engineers discussed their network configurations and provided a number of suggestions based on security and NAT devices present in the path of the transfer.
- **160 - American Museum of Natural History (AMNH):** During a routine NetSage demonstration for the American Museum of Natural History (AMNH), EPOC personnel noted possible poor performance of routine transfers between AMNH and the University of Arizona (UA). Examples of the suspected poor transfer performance were brought to the attention of engineers and administration at AMNH after the demonstration was complete.
- **161 - Rensselaer Polytechnic Institute (RPI):** The Chief Network Architect at NYSERNet contacted EPOC to discuss, and investigate, the performance impacts of configuring smaller than standard MTU settings on campus border routers at RPI to support the use of a commodity network Denial of Service (DoS) protection service. During the investigation, EPOC engineers discovered that the network hardware that is involved in delivering this service will transparently manipulate the size of network packets that cross the network segment independently of the end-devices communicating. The concerns raised by NYSERNET was that this could impact performance for certain large flows supporting science users, and may introduce other network behaviors (e.g., packet fragmentation or MTU blackholes). EPOC was not able to deliver conclusive results on the performance impacts, due to a lack of testing resources available on the RPI campus. Results from the prior work with NYSERNet and Syracuse University (Case #157), along with limited testing to a perfSONAR node on the RPI campus, indicated that modern network hardware and servers would be able to handle the MTU changes with limited impact to overall data transfer speeds.
- **162 - LEARN:** LEARN was awarded an NSF CC* grant (https://www.tx-learn.org/Grants/CC_Grant1925553.php) to work with their membership to create a regional infrastructure to support scientific research. LEARN staff requested time from EPOC to understand how the Deep Dive process could be integrated into their initial meetings with participants in this grant. EPOC gave a briefing to LEARN staff and some of the regional members, and provided materials that could be used virtually to conduct requirements analysis.
- **163 - Yale:** EPOC and Yale staff discussed the design best practices and tuning a Data Transfer Node (DTN), including feedback on processor selection, network architecture for a multi-10G-link setup, and RAM channel optimizations based on processor selection.
- **166 - National Oceanic and Atmospheric Administration (NOAA):** NOAA and EPOC staff met early in Quarter 1 to discuss the Deep Dive process and the potential for one to

be done virtually. NOAA identified a use case in acoustics as the basis for a Virtual Deep Dive and presentation with NOAA members. The NOAA Virtual Deep Dive is detailed in Section 5.

- **167 - LEARN:** LEARN staff procured several Ciena 3926 switches for use at various customer sites in the LEARN network. These switches are capable of running a virtualized instance of perfSONAR, and engineers were using this platform to do bandwidth tests from the LEARN core to the customers border router. Engineers noticed that at MTU 9000 they could achieve ~9Gbps speeds, while at MTU 1500 only ~3Gbps speeds were achievable. After consulting with EPOC engineers, it was suggested that LEARN engineers use the Linux-based tool mpstat to understand the load on the host processor during a bandwidth test at each MTU setting. This testing revealed that the host processor was running at 100% during bandwidth tests at MTU 1500 and was unable to keep up with the number of packets generated during a 10Gbps test. When the MTU was set to 9000, fewer packets were generated, which meant that the CPU load was smaller and the device achieved better performance. LEARN engineers, armed with a better understanding of the capabilities of the hardware, have now configured their test specifications to make the best available use of these resources.
- **168 - Southern Crossroads (SoX):** The President of SoX reached out to EPOC engineers for help troubleshooting data transfers from a newly installed DTN to several endpoints in the SoX network. EPOC engineers worked with him to perform perfSONAR tests to a number of remote nodes in the SoX regional network. The performance problem was isolated to the DTN device itself. Specifically, the DTN was running the Packetbeat software as one of its system logging tools. Testing with the Packetbeat software disabled showed that transfers to and from the DTN achieved expected performance.
- **169 - National Institute of Standards and Technology (NIST):** The supervisor of the NIST Boulder Network Operations Team reached out for EPOC recommendations related to Science DMZs and DTNs. They have a use case with data sets from 1-200TB and are about to make investments in their infrastructure. The EPOC team provided examples and answered questions.
- **170 - Lonestar Education And Research Network (LEARN):** During a recent round of network upgrades, engineers with the LEARN set up two new 10Gbps capable perfSONAR tester using Docker as an abstraction layer on top of dedicated physical hardware. Initial tests with these Docker instances indicated much lower than expected bandwidth, consistently lower than 1Gbps. LEARN engineers contacted EPOC for diagnostic and configuration assistance. EPOC engineers found that the TCP window size was not scaling from its initial size of ~8MB, which severely limited performance to any perfSONAR tester even a short distance from the regional network.
- **171 - Kennesaw State University:** The Director of Research Computing at Kennesaw State University reached out to discuss ways to increase utilization of research computing services by faculty and students. EPOC staff proposed doing a virtual Splash (lightweight Deep Dive) to help identify and document research use cases and issues seen by researchers, and to make recommendations based on that process. The Deep Dive was completed, and the report published in Quarter 3 [109].

- **172 - Indiana University (IU):** A network engineer from IU had a number of questions about low level transmission control protocol (TCP) behavior and interaction with various firewall deployments. EPOC staff discussed the issues and provided answers while encouraging the network engineer to test his theories in a lab environment.
- **173 - University of Arkansas (UArk):** The Trusted CI project had an engagement with University of Arkansas. UArk has a CC* award and are working towards deploying a number of Science DMZs at multiple campuses across the state connected by the Arkansas regional network (AREON). Trusted CI asked us to partner with them to document the best security practices and architectures for Science DMZs and DTN deployments. A joint paper addressing security and DMZ practices was written and published [118] in Quarter 4 along with a blog post announcing the publication [119]
- **174 - Kent State University (KSU):** The Chair of the Computer Science department at KSU reached out to EPOC to explain the best common practices for the operation of data transfer and computational devices that can be housed within a Science DMZ enclave. KSU was recently awarded a CC* computational grant, and was interested in trying to allocate computational jobs (e.g., provided by OSG software) on the same devices that were also tasked with data movement. EPOC staff explained some of the requirements for access control and policy on DTNs, and discouraged the use of the same machines for computation, but also put KSU in contact with a representative from Open Science Grid (OSG) to further explore possibilities.
- **175 - New York University - Abu Dhabi (NYU-AD):** EPOC was engaged by an engineer at NYU-AD who was testing Globus by attempting to read and write to an ESnet data transfer node (DTN). ESnet DTNs are only accessible via R&E networks and the NYU-AD routes were not being passed to ESnet by Internet2 due to policy. EPOC engineers explained how to work with Internet2 and ESnet to fix the problem.
- **176 - Lonestar Education and Research Network (LEARN):** LEARN staff reached out to discuss the EPOC Deep Dive process and how they can use it as an ongoing part of their CC* grant. EPOC staff answered some basic questions and are now tracking these possible Deep Dives, detailed in Section 9.
- **177 - University of South Dakota (USD):** EPOC continued providing assistance to USD as they prepared their CC proposal. Over a series of meetings, EPOC worked with USD and their regional partners to incorporate lessons learned from a previous EPOC virtual Deep Dive into their proposal. Specifically, EPOC advised that they integrate science stories to have a compelling CC regional narrative. EPOC provided a LoC for the proposal in Quarter 2 and has agreed to partner with USD if funded.
- **179 - Kent State:** Kent State staff requested an LoC from EPOC for their proposal to the CC* program. EPOC and Kent State staff discussed the proposal and reviewed the project summary. EPOC then provided a LoC and has agreed to partner if the proposal is funded.
- **180 - Lonestar Education And Research Network (LEARN):** In order to fulfill CC* grant requirements, LEARN engineers needed to provide shell access to perfSONAR testers installed at several campuses. LEARN engineers reached out to EPOC staff for guidance on best practices, citing security concerns with unfettered shell access. EPOC engineers responded with advice on how to construct a restricted secure shell

environment. LEARN took initial steps to implement this solution and will continue work on it in their own time, and did not need the ticket to remain open.

- **181 - I-Light:** I-Light staff reached out to EPOC for assistance in preparing a CC* proposal with their members Marian University and Ancilla College. Over a series of meetings, EPOC staff worked with the partners to help them better understand the solicitation, to identify science drivers, and to discuss possible areas for infrastructure improvements. At the end of the consultation, ILight decided the timing was not right for them to submit a proposal.
- **184 - Keystone Initiative for Network Based Education and Research (KINBER):** KINBER staff reached out to EPOC and requested a Letter of Collaboration for the planning grant proposal they submitted to the NSF CC* Solicitation. EPOC staff met with KINBER multiple times to discuss design review, hardware requirements, science drivers, and performance testing. EPOC provided the LoC and agreed to partner with KINBER if the proposal was successful.
- **186 - Ohio Academic Resources Network (OARnet):** EPOC staff consulted with OARnet engineers about peering with the Route Views project and how it can benefit OARnet and the R&E community. OARnet decided this was a lower priority at this time and will contact us if they need help in the future.
- **187 - University of North Carolina at Charlotte (UNCC):** EPOC has collaborated with the University of North Carolina Charlotte previously to provide a Letter of Collaboration for a CC* proposal, along with providing some details on how to secure science drivers, and propose CI design principles. This ticket focused on a follow-up to UNCC's award (NSF #2126116), which consisted of reviewing the network design, suggesting alternatives, and advising on implementation strategies.
- **188 - I-Light:** EPOC and I-Light engineers worked with Route Views engineers to configure peering between I-Light and the Route Views project.
- **189 - University of California San Francisco (UCSF):** A biology researcher at the University of California San Francisco (UCSF) contacted EPOC after having challenges when transferring large data sets to collaborators at Department of Energy sites. Initial investigations found that UCSF had multiple connections to CENIC (and thus access to ESnet peering points to reach DOE sites), but was having trouble accessing these via internal networking configurations. EPOC worked with UCSF and CENIC/PNWGP engineers to identify possible configuration changes, and worked to establish regular perfSONAR testing to ensure that the path and performance met expectations.
- **191 - Sun Corridor:** During a partner call, Sun Corridor staff expressed interest in peering with Route Views. EPOC staff introduced a Route Views engineer and sent information on why peering with Route Views is of benefit to the R&E community. Sun Corridor decided this was a lower priority at this time and will contact us if they need help in the future.
- **192 - NJEdge:** NJEdge staff reached out to EPOC to request a Letter of Collaboration for their CC* proposal. EPOC staff reviewed their proposal summary, provided a Letter of Collaboration, and agreed to collaborate and partner with NJEdge if their proposal was funded.

- **193 - Texas Association of Developing Colleges (TADC):** EPOC consulted with members and contractors representing the Texas Association of Developing Colleges (TADC), a consortium of private colleges in Texas, for assistance in preparing a CC* grant. The IT organization contracted to design the TADC regional DMZ requested assistance in choosing hardware for the network, but did not provide information on scientific use cases. Science drivers were ultimately not provided and EPOC was unable to recommend hardware without that information.
- **194 - University of Arkansas (UArk):** The PI for the NSF SHARP-CCI program reached out to EPOC for support. A talk was given at their kickoff meeting [107], focusing on the EPOC activities that could support SHARP-CCI members and what makes up a strong CI Plan for an institution. The SHARP PI will reach out for additional support as needed.
- **196 - American Museum of Natural History (AMNH):** EPOC started a conversation with AMNH after the virtual CC* PI meeting in 2021 to follow up on a request to get more information on Deep Dives and hosting possible future training activities. AMNH staff let us know they would reach out in the future as needed.
- **197 - Massachusetts Green High Performance Computing Center (MGHPCC):** EPOC started a conversation with MGHPCC after the CC* PI meeting in 2021 to follow up on a request to provide more information on the Deep Dive process EPOC uses to identify and support technology for scientific use cases.
- **199 - Yale:** EPOC staff started a conversation with members of the Yale University High Performance Computing group after the virtual CC* PI meeting in 2021 to follow up on a request to get more information on Science Deep Dives and NetSage monitoring. After some discussion, it was found that the Yale campus uses DHCP for IP address assignments, so IP addresses may change frequently. This will make the deployment of a campus NetSage deployment challenging. Yale continues to discuss a possible Deep Dive with EPOC, but has had staffing shortages which have prevented this from occurring.
- **200 - Memorial Sloan Kettering Cancer Center (MSKCC):** EPOC was contacted by NYSErNet to assist MSKCC on a possible proposal to either CC* or CICI on the topic of creating a testbed to evaluate medical device security. Due to the time between the request and the solicitation deadlines, this was pushed to a future solicitation. The ticket was closed after MSKCC indicated they had enough information to proceed at a later date.
- **201 - Claflin University:** EPOC was contacted by staff at Claflin University, an HBCU, on participation in the Data Mobility Exhibition (DME) activity along with assistance on best common practices for designing and configuring data transfer devices. EPOC has provided some information about available resources, including documentation and video content.
- **203 - Louisiana Optical Network Initiative (LONI):** LONI started an ongoing conversation with EPOC based on a YouTube posted talk on performance tuning. Questions were based around new TCP algorithms and performance tuning of a circuit.
- **205 - Kennesaw:** Kennesaw personnel requested help implementing a perfSONAR solution to understand their baseline network performance and identify potential transfer

performance issues. EPOC responded with best practices for implementing perfSONAR and offered the use of a loaner 10G perfSONAR node which was declined.

- **206 - Thomas Jefferson National Accelerator Facility (JLab):** During routine evaluation of data performance at Department of Energy sites, EPOC contacted staff at Jefferson Laboratory (JLab) to help understand lower than expected values for data transfer to some collaborators. Working with JLab engineers, EPOC was able to identify a misconfigured Equal Cost MultiPath (ECMP) connection to ESnet within the JLab infrastructure, due to the way that JLab configures and load balances their science and enterprise traffic flows across two ESnet connections. The ECMP configuration was corrected.
- **207 - University of Nevada Reno (UNR):** A researcher from University Nevada Reno asked for an LoC for the NSF CSSI solicitation (#20592). The request was for EPOC to test a researcher's data movement software and to recommend it during EPOC consultations. The researcher was informed that EPOC doesn't do beta-software testing and only recommends production software with support teams, but that we could support them by offering a venue to speak at. The request was withdrawn.
- **208 - University of Hawaii:** Staff at the National Radio Astronomy Observatory (NRAO), in collaboration with staff at the University of Hawaii, reported that perfSONAR devices were exhibiting larger than normal traffic patterns. Speculation was that a recent CVE issued on software affiliated with the perfSONAR toolkit was actively being exploited. EPOC, NRA, UH, and the perfSONAR project collaborated on ways to verify configurations, as well as possible mitigations and protections that can be put in place to ensure the machine in question was checked, validated, and possibly reconfigured to prevent abuse in the future.
- **209 - Indiana University (IU):** A researcher from Indiana University Bloomington asked for an LoC for the NSF CSSI solicitation (#20592). The request was for EPOC to recommend a prototype for software during EPOC consultations. The researcher was informed that EPOC only recommends production software with support teams, but that we could support them by offering a venue to speak at. The request was withdrawn.
- **211 - University of Hawaii:** EPOC was contacted by the University of Hawaii (UH) regarding the best process to transfer larger amounts of data out of CERN for a research project. After an initial investigation, it was found that while the University had adequate DTN resources (hardware and software), the scientific activity in question used a different set of tools. UH IT staff installed the proper software tools onto their existing DTN hardware resources to facilitate the scientific portion of the data movement exercise. In parallel, UH IT continued to work with EPOC to test and validate the underlying networking path with perfSONAR and Globus using ESnet-affiliated resources at CERN. The performance baselines met expectations.
- **215 - University of Cincinnati (UC):** The Associate Director of Research Computing and Data at University of Cincinnati has reached out with a request to review and help with their upcoming Science DMZ upgrade. EPOC staff provided pointers and advice for UC engineers to review. The engineers reported that the information provided gave them the answers they needed.

- **218 - Duquesne:** Duquesne personnel were getting redirected to an Internet2 web address when trying to download the current perfSONAR image from the perfsonar.net website. This error turned out to be related to their DNS configuration and was corrected by on-site engineers.
- **219 - University at Buffalo (UB):** EPOC was contacted by the UB High Performance Computing Group for assistance with defining security policy to support data movement and manage internal and external users. EPOC did not have sufficient subject matter expertise with respect to security policy for supporting HPC users and worked with staff at NERSC to advise on best practices used at large computing facilities, and how they may be implemented by the University at Buffalo.
- **221 - University of Buffalo (UB):** EPOC was contacted by the UB High Performance Computing Group for assistance in locating software that can be used to monitor and map their network infrastructure. EPOC convened a call that discussed various open source and commercial options in this space.
- **224 - Howard Hughes Medical Institute (HHMI):** EPOC was contacted by the Scientific Computing Systems at HHMI Janelia Research Campus for assistance in testing to ESnet DME test points. After an investigation, it was found that HHMI was not connected to the R&E networking infrastructure, thus they cannot reach ESnet DME test points due to policy. EPOC worked with HHMI to test against other public DME test points, and started a conversation with the Mid-Atlantic Crossroads (MAX), the regional provider in Maryland and Virginia, about ways HHMI may connect in the future.
- **226 - University of Michigan (UM):** Staff from UM contacted EPOC personnel to schedule a refresher course in CI engineering for new staff. EPOC scheduled a call, and went over various aspects of Science DMZ design, while evaluating the current design of the UM research network.
- **230 - Alabama A&M University (AAMU):** Staff at AAMU contacted EPOC for information on training resources. EPOC provided links to recent virtual events, and added members to an events announcement list for future opportunities.
- **231 - Abbvie:** Staff from Globus contacted EPOC members for assistance in explaining CI engineering approaches to Abbvie, a commercial entity that performs research with National Labs and universities. EPOC conducted a call to explain aspects of the Science DMZ design and offered to be available to answer questions in the future.
- **234 - Heinrich Heine Universität Düsseldorf (HHUD):** A researcher at HHUD contacted EPOC after they encountered issues in testing the ESnet DME test points. As with prior reports, the problem was found to be that the University was not connected to an R&E network. After working with the researcher and DFN, the NREN for Germany, re-testing was successful.

Ongoing consultations included:

- **142 - University of Alaska Fairbanks (UAF):** EPOC staff were contacted by staff at UAF, and researchers at the International Arctic Research Center, to assist in specifying networking connectivity options that may match scientific drivers, and fit within budget and technology capabilities to their facility. Initial discussions involved connectivity to Fairbanks Alaska, and the realities of working with the limited and expensive

telecommunications resources. EPOC advised on network design to support science, including Science DMZ best practices, data transfer nodes, and the need for deep buffered switches to support long distance data transfers. Conversations also discussed details about applying to the CC* program as a possible source of funding for some of the upgrades. UAF did not succeed in getting a proposal together for the 2021 or 2022 CC* program deadlines, but is still interested in pursuing one in 2023 when it is released.

- **178 - Saint Louis University (SLU):** The newly appointed Vice President of Research (VPR) at SLU has asked for EPOCs help identifying the network needs of three science use cases. EPOC staff are working with SLU on virtual Splash (lightweight Deep Dive). In parallel, EPOC and GPN staff plan to investigate network issues experienced by some of the researchers on campus. In Quarter 3, EPOC sent SLU a self contained data transfer node (DTN) for iperf and Globus testing. During Quarter 4, EPOC Staff engaged the GPN CyberTeam to help with the consultation. The CyberTeam provided a perfSONAR node and are working with SLU engineers, security team, and management to get approval for installation. While waiting for the approvals SLU engineers used the EPOC provided node to start initial testing using iPerf. Testing at this time has only been from behind the SLU firewall to iPerf nodes installed in the GPN regional network. In Year 5, the CyberTeam plans to finish testing and work with EPOC to make network upgrade recommendations.
- **182 - Keystone Initiative for Network Based Education and Research (KINBER):** During a routine EPOC partner meeting, staff from the Keystone Initiative for Network Based Education and Research (KINBER) requested assistance with their NetSage installation as they were about to change their router configurations at their two major POPs. They indicated that they would be unable to send flow or SNMP data to NetSage for a period of time and would reach out when their configuration changes were complete. At this time there is no indication of when flow and SNMP collection will resume.
- **183 - Keystone Initiative for Network Based Education and Research (KINBER):** During a routine EPOC partner meeting, staff from the KINBER requested assistance with an upcoming project to improve connectivity to several public library sites within their jurisdiction, specifically in the area of network performance evaluation before and after the network upgrades. Initially EPOC personnel suggested that perfSONAR nodes be installed at each site, and that they be incorporated into a testing mesh, with periodic tests run against nodes at KINBER's POPs in Philadelphia and Pittsburgh. While discussion is continuing about this plan, other alternatives are being explored due to lack of staffing at individual library sites.
- **185 - Lonestar Education And Research Network (LEARN):** EPOC staff are consulting with LEARN engineers and the Route Views project staff to get LEARN routers peered with the Route Views collectors.
- **190 - Sun Corridor:** During a partner call, Sun Corridor staff expressed interest in updating and expanding their perfSONAR testing infrastructure. This work is currently on hold while Sun Corridor staff address other priorities.
- **195 - University of Connecticut (UConn):** EPOC engaged with a researcher from the University of Connecticut after the virtual CC* meeting in 2021 on a long-standing data

transfer issue between his facility and Jefferson Lab (JLab), a Department of Energy facility that hosts the GlueX nuclear physics experiment. The initial report stated that the researcher was experiencing slow data transfers in trying to move data to use some experimental analysis software on the Open Science Grid (OSG). EPOC investigations with JLab revealed several layers of possible performance problems, including: an incorrectly configured network at JLab (see ticket #206), a need to increase capacity between JLab and ESnet to 100Gbps, an older Science DMZ that needed upgrades for better performance, , and older data movement hardware in the path that was a known performance bottleneck. EPOC continues to engage with all parties as upgrades are performed. Supply chain issues have slowed this ticket considerably.

- **198 - University of Colorado - Colorado Springs (UCCS):** EPOC staff had a request from UCCS staff for additional information about how Deep Dives are performed and if the process would work for UCCS. This request is ongoing.
- **202 - University of Arkansas (UArk):** EPOC staff are working with staff from UArk to complete some Science DMZ network performance baselining using perfSONAR tests as well as understanding the performance baseline of their HPC DTNs using Data Mobility Exhibition datasets. Initial results were inconsistent and showed poor performance. The poor performance illustrated different characteristics of common configuration tuning issues. UArk engineers adjusted several network settings, but the performance did not noticeably change. EPOC staff learned that the UArk DTN was a virtual machine, and suggested that instead the Globus software should be installed directly on the server for better results. This engagement is ongoing.
- **204 - Ocean State Higher Education Economic Development and Administrative Network (OSHEAN):** Poor transfer performance was observed on the OSHEAN regional network, most notably from Brown University to several European endpoints, by EPOC personnel using NetSage. Initial investigation revealed an MTU mismatch between Internet2 and the OSHEAN, and initial steps toward resolving this problem were taken by OSHEAN engineers. After further discussion, a ticket has been opened with the Indiana University GlobalNOC, which manages the OSHEAN network, and a maintenance window has been scheduled to correct the MTU mismatch.
- **210 - The South African National Research and Education Network (SANReN):** SANReN reached out to EPOC with a request from astronomy researchers to better understand performance of data flows between the MeerKAT radio telescope in South Africa and the NRAO data archives at the University of Virginia. After a series of meetings, EPOC staff are now working with SANReN to expand the current SANReN NetSage deployment to include flow data for both primary and secondary transatlantic paths and to tag relevant flows in the NetSage Science Registry.
- **212 - Cloud Enhanced Open Software Defined Mobile Wireless Testbed for City-Scale Deployment (COSMOS):** Researchers from the COSMOS project (NSF #1827923) reached out to EPOC engineers to discuss ways to increase transfer speeds between the City College of New York (CCNY) and Kyutech Institute of Technology in Japan. After initial discussion, it was decided that the best path forward would be to create 3 VLAN's directly between CCNY and Kyutech. EPOC engineers started an email conversation with relevant parties along the path including TransPAC, Pacific

Wave/PNWGP, JGN, Internet2, and NYSERNet and VLAN configuration has begun. We expect it to be complete and ready for testing and production early in Year 5.

- **213 - Purdue University:** Reached out to Purdue to follow-up on a number of Data Mobility tests they completed. Purdue is currently testing multiple HPC DTNs and connected storage systems for their HPC infrastructure. We hope to work with Purdue to test their NSF ACCESS Anvil Super Computer DTNs
- **214 - University of Alabama Birmingham (UAB):** The DME program reviewed their logs and reached out to UAB staff after finding out that they had initiated thousands of tests in a short period of time. A student had used a jupyter notebook to create a process to perform automated testing to and from their campus to research how much the institutional firewall affected the measurement results. EPOC discussed their approaches to this testing and reviewed the results. UAB will be revising their network and will be testing DME again as well as Upgrading their DTNs to Globus Connect Servers (GCS) version 5..
- **216 - St. Jude Children's Hospital Network (St. Jude):** A consultant that is assisting St. Jude contacted EPOC for best common practices for Science DMZ Design. After providing some offline documentation, St. Jude has requested a follow-up call to discuss their use case, and possible design strategies they can follow. This request is ongoing.
- **217 - Baylor University:** Staff at Baylor University were contacted by EPOC to install and evaluate a "container"-based instance of EPOC's Modern Research Data Portal (MRDP). The goal of this interaction is to get feedback on deployment and documentation, and to provide a service to a Baylor Researcher profiled in the 2020 Baylor Deep Dive that could benefit from the software assistance. This request is ongoing.
- **220 - National Radio Astronomy Observatory (NRAO):** NRAO staff reached out for assistance regarding the configuration of several new perfSONAR nodes. These nodes had replaced older nodes, and staff had attempted to import the configurations from the old nodes. While the testing configurations were successfully imported, the data from these tests wasn't being stored on the nodes. EPOC personnel found several configuration errors and corrected them, however test results are still not being stored. NRAO staff are currently busy with other priorities, but once their schedule allows this work to correct the node configuration will resume.
- **222 - Colorado State University (CSU):** Staff at CSU reached out to EPOC for assurance in redesigning their Storage Area Network (SAN), and ways that data movement hardware and software could be integrated to assist with several research drivers. EPOC scheduled a call, and walked through several CI engineering topics (networking, data transfers, security). The request is ongoing as we wait for additional implementation questions.
- **223 - ESnet:** Globus contacted ESnet to assist with tickets that had been opened by their users regarding testing to ESnet DME test points. As with prior tickets, a brief EPOC investigation found that the end-sites in question were not connected to R&E networks and policy prevented the tests from being completed. This ticket remains open, pending the delivery and installation of new ESnet DME test points in 2022 with the DME datasets. Other DME test points were shared to support non-R&E test data transfers.

- **225 - Lafayette College:** Lafayette College contacted EPOC for assistance on general CI capabilities for the college and how they could better support several science use cases. The college is partnering with KINBER on a CC* regional solicitation, but has not heard back on the outcome as of reporting. EPOC has provided some insight, and will continue to work with Lafayette into 2022.
- **227 - University of Texas Health San Antonio (UTHSA):** UTHSA staff reached out to members of EPOC after initially engaging with the Texas Advanced Computing Center (TACC) on best practices for data transfers. Initial conversations found a number of areas that UTHSA could improve upon, which included normalizing their routed networking to prefer R&E connectivity delivered by the UT system and LEARN, considering the installation of a Science DMZ infrastructure, using perfSONAR for performance monitoring, and installing a dedicated data movement hardware and software solution based on Globus. This engagement is ongoing.
- **228 - University of Central Florida (UCF) -** UCF staff reached out to EPOC reporting that an astronomy researcher on their campus was having issues transferring data from the Evryscope project in Chile to several sites in the US. We are in the process of setting up a call with stakeholders to learn more about the issues they are experiencing.
- **229 - University of Utah (UU):** Globus reached out to EPOC for assistance in helping a researcher working on the Sloan Digital Sky Survey (SDSS) project migrate data from resources at the UU to several sites in Europe. EPOC helped to connect Globus and IT staff at UU to better understand the networking path. Initial investigations have shown that the UU and (Utah Education and Telehealth Network (UETN) connectivity is sufficient, which indicates that problems may exist outside of the United States. EPOC is keeping this ticket open, pending notification from Globus.
- **232 - Smithsonian Data Science Lab (SDSL):** EPOC was contacted by the members of the SDSL for additional information about best common practices for data transfers. EPOC provided some documentation on data transfer hardware, software, and networking design. This engagement is being kept open, pending some internal discussions at SI regarding next steps.
- **233 - Oak Ridge National Laboratory (ORNL):** EPOC is engaging with the Atmospheric Radiation Measurement Facility (ARM), located at ORNL, to seek help in connecting remote sensor sites around the world. ARM uses a small array of atmospheric and weather sensors to collect data, and then sends the results back to their main data facility at ORNL. Due to the location of these sensors, they are often in areas that do not have strong terrestrial connectivity options and must rely on cellular or satellite connections. EPOC and ESnet are engaged with ARM, and will be working with them to identify locations and connectivity options with and without R&E partners.

4.C Metrics

Table 1: A summary of Roadside Assistance and Consultation Cases in Year 4. Green rows are completed cases.

ID	Main Site	EPOC Partner	Type	Start Date	End Date	Area of request	Asked by: Eng, Scientist, other	Science Domain	Primarily R(ch), E(du), O(ther), (N)etwork	Size : S, M, L
63	NOAA	FRGP	Cons	12/16/19	4/21/21	DD	Eng	Infra	O	-
76	NCAR	FRGP	RA	3/6/20	6/17/21	Trans Perf	Eng	Climate	R	L
105	LC	KINBER	Cons	8/13/20	6/30/21	Grants, DMZ, DTN	Eng	Infra	E	S
107	Arecibo	TACC, Globus	RA	8/21/20	6/1/21	Trans Perf	Eng	Astro	R	L
113	TAMU	TACC, LEARN	Cons	9/10/20	6/30/21	Perf	Eng	Infra	R	L/M
117	AI	PNWGP	Cons	10/21/20	9/2/21	Research	Eng	Infra	R	L
121	EHT	NEAAR/TP/ UH Astro, PIREN	RA	10/28/20		Trans Perf	Sci	Astro	R	L
135	TACC	TACC	Cons	1/5/21	6/20/21	Arch	Other	Infra	E	L
142	UAF	PNWGP	Cons	1/28/21		Grants, Arch	Eng	Infra	R / E	M
143	ASU	SCN	Cons	2/9/21	8/4/21	Arch	Eng	Infra / Bio	R / E	L
147	SARAO	Globus	Cons	2/16/21	8/18/21	Transfer Perf, Globus	Eng	Infra	R	L
152	Brown		Cons	2/24/21	5/3/21	DMZ, Firewall	Eng	Infra	R / E	M
153	SLU	GPN	Cons	2/24/21	5/26/21	Arch	Eng	Infra	E	M
154	NYSERNet		Cons	3/4/21	4/5/21	NS	Eng	Infra	R	-
155	TAMUSA	LEARN	Cons	3/8/21	6/15/21	DMZ, Research	Eng	Infra	E	S
157	Syracuse		Cons	3/25/21	5/11/21	Perf, Security	Eng	Infra	E	M
158	NYULH		Cons	3/26/21	5/14/21	DMZ	Eng	Medical	R	M / L
159	Yale	Globus	Cons	3/31/21	6/10/21	Trans Perf	Eng	infra	R / E	L
160	AMNH	SCN	Cons	3/31/21	9/29/21	Trans Perf	Eng	Infra	R/E	-
161	RPI		Cons	3/31/21	5/11/21	Trans Perf, Security	Eng	infra	R / E	M
162	LEARN	LEARN	Cons	4/7/21	5/3/21	DD	Eng	infra	R / E	L
163	Yale		Cons	4/13/21	5/12/21	DTN	Eng	infra	R / E	L
164	KNU	TP	RA	4/15/21	5/16/21	Trans Perf, Routing	Eng	infra	R	S
165	NII	TP	RA	4/15/21	5/16/21	Trans Perf, Routing	Eng	infra	R / E	S

ID	Main Site	EPOC Partner	Type	Start Date	End Date	Area of request	Asked by: Eng, Scientist, other	Science Domain	Primarily R(ch), E(du), O(ther), (N)etwork	Size : S, M, L
166	NOAA	FRGP	Cons	4/23/21	6/16/21	DD	Eng	environmental	R	L
167	LEARN	LEARN	Cons	5/13/21	6/30/21	Trans Perf, PS	Eng	Infra	R / E	-
168	SoX	SoX	Cons	6/7/21	6/30/21	TransPerf, DTN	Eng	infra	O	-
169	NIST	FRGP	Cons	6/10/21	8/4/21	DMZ, DTN	Eng	TBD	R	L
170	LEARN	LEARN	Cons	6/21/21	9/29/21	PS	Eng	infra	R	S
171	KSU	SoX	Cons	7/8/21	11/3/21	DD	Other	infra	E	M
172	IU	iLight	Cons	7/13/21	7/28/21	Security, Trans Perf	Eng	infra	R	L
173	UArk	TrustedCI	Cons	7/20/21	1/11/22	DMZ, Security	Eng	infra	E	M
174	KSU	OARnet	Cons	8/5/21	8/19/21	DMZ, Security	Eng	infra	E	M
175	NYU	NEAAR, I2	Cons	8/17/21	8/30/21	DTN, Routing, Globus	Eng	infra	E	L
176	LEARN	LEARN	Cons	8/18/21	8/23/21	DD	Eng	Infra	O	-
177	S. Dakota	GPN	Cons	8/18/21	10/12/21	Grants, DD	Eng	Infra	E	M
178	SLU	GPN	Cons	8/19/21		Trans Perf, DMZ	Other	infra	E	M
179	KSU	OARnet	Cons	8/19/21	10/12/21	Grants	Other	infra	E	L
180	LEARN	LEARN	Cons	8/19/21	10/25/21	Security, PS	Eng	infra	R, E	M
181	iLight	iLight	Cons	8/19/21	9/29/21	Grants	Eng	Infra	N	S
182	KINBER	KINBER	Cons	8/20/21		NS	Eng	infra	N	S
183	KINBER	KINBER	Cons	8/20/21		PS	Eng	infra	N	S
184	KINBER	KINBER	Cons	8/20/21	10/12/21	Arch, Grants	Eng	Infra	N	S
185	LEARN	LEARN	Cons	8/23/21		Routing	Eng	Infra	N	-
186	OARnet	OARNet	Cons	8/23/21	11/30/21	Routing	Eng	Infra	N	-
187	UNCC		Cons	8/23/21	1/7/22	Arch	eng	infra	R, E	M
188	I-Light	I-Light	Cons	8/24/21	9/13/21	Routing	Eng	Infra	N	-
189	UCSF	PNWGP	Cons	8/25/21	1/18/22	Trans Perf, Routing	Scientist	Bio	R	M
190	Sun Corridor	SCN	Cons	8/26/21		PS	Eng	Infra	N	-
191	Sun Corridor	SCN	Cons	8/26/21	12/13/21	Routing	Eng	Infra	N	-
192	NJEdge		Cons	9/2/21	9/17/21	Grants	Eng	Infra	N	-
193	TADC	LEARN	Cons	9/2/21	10/12/21	Grants, DMZ	Eng	Infra	O	-

ID	Main Site	EPOC Partner	Type	Start Date	End Date	Area of request	Asked by: Eng, Scientist, other	Science Domain	Primarily R(ch), E(du), O(ther), (N)etwork	Size : S, M, L
194	Uark	GPN, Trusted CI	Cons	9/7/21	10/29/21	CI Plans	Eng	Infra	R, E	M, L
195	UConn		Cons	9/21/21		Trans Perf, DME	Scientist	Physics	R, E	M
196	AMNH		Cons	9/21/21	1/18/22	DD, Training	Eng	Infra	R	M
197	MGHPCC		Cons	9/21/21	1/4/22	DD	Eng	Infra	R, E	M, L
198	UCCS	FRGP	Cons	9/21/21		DD	Eng	Infra	R, E	S
199	Yale		Cons	9/21/21	2/23/22	DD, NetSage	Eng	Infra	R, E	L
200	MSKCC		Cons	9/27/21	10/19/21	Grants	Eng	Medical	R	M, L
201	Claffin		Cons	11/4/21	1/26/22	DME, Arch	Eng	Infra	E	S
202	Arkansas	GPN	Cons	11/12/21		DME	Other	Infra	E	M
203	LONI		Cons	11/15/21	1/12/22	Trans Perf	Eng	Infra	N	M
204	OSHEAN	NEA3R	Cons	11/17/21		Trans Perf	Eng	Infra	N	M
205	Kennesaw		Cons	11/18/21	3/9/22	Trans Perf, PS	other (VPR)	Infra	E	M
206	JLab		Cons	11/18/21	11/18/21	Trans Perf, Routing	Eng	Nuclear Physics	R	M
207	UNR		Cons	11/29/21	12/9/21	Grants	Scientist	CS		
208	UH	PIREN, UH Astro	Cons	11/29/21	1/11/21	PS, Security	Eng	Infra	R, E	L
209	IU	I-Light	Cons	12/2/21	12/9/21	Grants	Scientist	CS	R, E	
210	SanReN	NEA3R	Cons	12/3/21		Trans Perf, NS	Eng	Astro	R	L
211	UHawaii	PIREN	Cons	12/6/21	1/31/22	DTN, Trans Perf	Eng	Physics	R, E	M
212	COSMOS	TP, PNWGP, I2	Cons	12/9/21		Arch, Routing, Trans Perf	Scientist	IOT	R	M
213	Purdue	iLight	Cons	12/13/21		DME, Trans Perf	Eng	Infra	R, E	L
214	UAB	SoX	Cons	12/14/21		DME, Trans Perf	Eng	Infra		
215	UConn		Cons	1/12/22	2/9/22	DMZ	Other	Infra	E	L
216	St Jude		Cons	1/12/22		DMZ, Arch	Eng	Medical	R	M
217	Baylor	LEARN	Cons	1/14/22		DME, Trans Perf	Eng	Infra	R & E	M
218	Duquensne	KINBER	Cons	1/19/22	1/24/22	PS	Eng	Infra	E	S
219	UB		Cons	1/25/22	1/25/22	Security	Eng	Infra	R & E	M
220	NRAO		Cons	2/3/22		PS	Eng	Astro	R	L
221	UB		Cons	2/3/22	2/7/22	Arch	Eng	Infra	R & E	M

ID	Main Site	EPOC Partner	Type	Start Date	End Date	Area of request	Asked by: Eng, Scientist, other	Science Domain	Primarily R(ch), E(du), O(ther), (N)etwork	Size : S, M, L
222	CSU	FRGP	Cons	2/4/22		Arch, DME	Eng	Infra	R, E	M
223	ESnet	Globus	Cons	2/14/22		Arch, DME	Eng	Infra	R	L
224	HHMI		Cons	2/16/22	3/10/22	DME, DTN	Eng	Medical	R	L
225	Lafayette	KINBER	Cons	2/23/22		DMZ, Arch	Eng	Infra	E	S
226	UM		Cons	2/24/22	3/17/2022	DMZ, Arch, Training	Eng	Infra	R&E	L
227	UTHSA	LEARN, TACC	Cons	2/24/22		DTN, Trans Perf, DME	Eng	Medical	R	L
228	UCF		RA	3/2/22		Trans Perf	Other	Astro	R	M
229	UU	Globus	Cons	3/4/22		Trans Perf	Eng	Astro	R	L
230	AAMU		Cons	3/4/22	3/4/2022	Training	Eng	Eng	E	S
231	Abbvie	Globus	Cons	3/10/22	3/10/22	DMZ, Globus	Eng	Pharma /Med	R	L
232	SDSL		Cons	3/23/22		DME	Scientist	Data Science	R	L
233	ORNL	SoX	Cons	3/23/22		Arch, DMZ	Eng	Climate	R	L
234	HHUD	Globus	Cons	3/14/22	3/29/22	DTN, Routing	Eng	Infra	E	M

4.D Additional Activities Related to Roadside Assistance

Over the course of Year 5, we gave 9 presentations which included material on Roadside assistance [92, 102, 103, 105, 107, 112, 122, 123, 124]. In addition, based on the issues that Roadside Assistance and Consultation Cases that we have addressed most frequently, we have created or are working on following Best Practice documents:

- It's All About the MTU's [76], which describes why the Maximum Transmission Unit (MTU) configuration parameters matter for end to end network performance and how to troubleshoot MTU issues.
- Affordable 10G DTN build overview using the Dell server configurator as a video [94] and a white paper [116].
- A guide to understanding why a 100G connected perfSONAR or DTN node may not always achieve 100G transfer rates, which is expected to be completed in Year 5 Quarter 1.
- A Science DMZ starter pack, which will consist of several documents giving the background, how to find and evaluate DMZ use cases, how to select a 10G DTN, and what the security aspects are. This will be a focus of Year 5.

EPOC is also working with our regional partners to document known good and well connected perfSONAR nodes in their networks. These nodes will be used during Roadside Assistance by EPOC engineers to troubleshoot reported network issues.

As a side project, we also reached out to a set of past CC* awardees in email to offer help with design review, including assistance with planning, engaging with resources, and recommendations to help complete their funded goals. Of the 15 participants contacted between July and September 2021, we did not receive a single response, so in Year 5 we are re-evaluating this approach.

5. Deep Dives

Deep Dives aim to understand the full research pipeline for collaborative teams and suggest alternative approaches for the scientists, local CI support, and national networking partners as relevant to achieve the long-term research goals via workflow analysis, storage and computational tuning, and identification of network bottlenecks. We have adapted the ESnet facilities approach for work with individual science groups, which is documented at: <https://epoc.global/wp-content/uploads/Application-Deep-Dive-Description.pdf>. Jason Zurawski is the lead for this area.

5.A Shift to Virtual, Streamlined Deep Dives (“Splash”)

As COVID-related travel restrictions were put in place, EPOC shifted to Deep Dives that could be executed completely virtually. The Virtual Deep Dive approach, sometimes called a Splash, involves using more streamlined Case Study google documents with additional video sessions to “train the trainers”. The online sessions assist the local IT staff to understand the Deep Dive structured conversation approach so they can work with the individual science groups to fill out the application Case Studies. Once both the Application and Technical Case Studies are collected, we schedule a series of Focus Group video calls to walk through the data and try to identify the CI needs and requirements. The full set of participants then update the findings and we combine the data and observations from the focus groups into a report that is reviewed and then made public.

In Year 4 Quarter 4, some travel restrictions were lifted, but we are still using this technique in many cases and plan to continue to offer it as an alternative.

5.B In Progress Deep Dives

The following Deep Dives activities are in progress:

- **Arizona State University / Sun Corridor Network - Hybrid:** In August, 2019, Arizona State University reached out to EPOC to host a potential Deep Dive of campus and regional requirements. An in-person event was planned for January 2022, but this

switched to virtual and took place January 10-12. A report is in progress which we hope to publish in early Year 5

- **LEARN/South Plains College- Virtual:** LEARN is working with several Texas community colleges to start a regional Deep Dive program to execute a CC* grant. EPOC is assisting LEARN and the first college, South Plains College, to gather data and create a report virtually. A virtual meeting was conducted in November 2021 with LEARN and South Plains, and a report is expected in mid 2022.
- **OARnet/Sinclair College - Virtual:** OARnet is working with several Ohio community colleges to start a regional Deep Dive program to execute a CC* grant. EPOC is assisting OARnet and the first college, Sinclair College, to gather data and create a report virtually. A virtual meeting was held in December 2021 with OARnet and Sinclair staff, and a report is expected in mid 2022. Other participants in the grant may follow the same pattern or work to have in-person events in 2022.
- **St. Louis University (SLU) - Virtual:** SLU had filed a Roadside Assistance request to help with network architecture, but this turned into performing a limited-reach Deep Dive with EPOC and GPN in January [53] to understand use cases. EPOC is working to publish these findings in Year 5.
- **LEARN/McLennan Community College:** LEARN is working with several Texas community colleges to start a regional Deep Dive program to execute a CC* grant. An in-person meeting was conducted in March 2022 with LEARN and McLennan [59], and a report is expected in mid 2022.
- **LEARN/Midland College:** LEARN is working with several Texas community colleges to start a regional Deep Dive program to execute a CC* grant. An in-person meeting is being planned for April 2022 with LEARN and Midland, and a report is expected in mid 2022.
- **San Diego State University Deep Dive:** EPOC is partnering with San Diego State University to perform a Deep Dive of campus drivers. The goal is to perform this as a hybrid event, with virtual questions that started in February of 2022, and an in-person event scheduled for April of 2022. A report is expected in Summer of 2022.

5.C Completed Deep Dives

The following Deep Dives activities were completed in Year 4:

- **NOAA Marine Mammal Acoustics (NOAA NMFS) - Virtual:** Staff from NOAA N-WAVE approached EPOC to profile a NOAA science area and to learn about the Deep Dive process so it can be adapted to other NOAA use cases. Virtual meetings with researchers and IT staff from NOAA were held in summer 2021, and a final presentation [88] occurred at NOAA N-Wave Joint Engineering and Technical Interchange (JETI) conference [23]. The effort was completed with a final report in August 2021 [91].
- **University of Central Florida (UCF) - Virtual:** Staff from UCF approached EPOC to stage a Deep Dive for the campus. Due to COVID-related travel restrictions, most of the activities to discover and discuss the scientific drivers were done virtually in late 2020 and early 2021. The goal was to produce content that can be used to justify a CC*

proposal for campus in 2021 or 2022. The final report was completed in August 2021 [93].

- **Kennesaw State University (KSU) - Virtual:** KSU had filed a Roadside Assistance request to help with network architecture but performed a limited-reach Deep Dive with EPOC to understand use cases. EPOC published this report in November 2021 [109].
- **University of South Dakota (USD) - Virtual:** Staff from USD approached EPOC to stage a Deep Dive for the region, pulling in participants from other local institutions, including South Dakota State, Black Hills State, and Northern State. Due to COVID-related travel restrictions, most of the activities to discover and discuss the scientific drivers were done virtually in late 2020 and early 2021. Virtual meetings with researchers and IT staff from the region were held in the Spring of 2021. EPOC published this report in October 2021 [108].

5.D Abandoned Deep Dives

The following Deep Dives were abandoned due to lack of interest and activity:

- **Oregon State University:** In April 2019, members of Oregon State University contacted EPOC staff about a possible EPOC Deep Dive to profile their campus research and the regional network for the state, LinkOregon. Dates and focus areas were discussed but have been canceled due to COVID-19.
- **Texas A&M University San Antonio (TAMUSA):** Staff from TAMUSA approached EPOC regarding Science DMZ design help but were challenged in designing a network that fit scientific requirements. EPOC, LEARN, and TAMUSA discussed a Virtual Deep Dive to gather scientific use cases, but the activity was abandoned due to lack of TAMUSA staff interest to work via virtual mechanisms.

5.E Metrics

Table 2: Metrics for Deep Dive activities.

Meet Date	Appl name	Public/Private	Audience	Offered or Req	Head Count	Issues Identified	Complete Date
Jan 2022	Arizona State Univ, Sun Corridor	Priv	ASU, UofAZ, NAU, and Sun Corridor Network staff	Req	TBD	Storage, compute, and general need for CI support	Expected Summer 2022
Summer, Fall 2021	Kennesaw State University (KSU)	Priv	KSU Researchers & Staff	Offered	10 Virtual	Connectivity, HPC availability	Published November 2021
Jan 2022	LEARN / South Plains College	Priv	LEARN CC* Grant Participants	Req	8 Virtual	Software fees, lack of storage	Expected Spring 2022
Jan 2022	OARnet / Sinclair College	Priv	OARnet CC* Participants	Req	20 Virtual	Wireless testbeds, lack of storage	Expected Summer 2022
Summer,	St. Louis	Priv	SLU Researchers &	Offered	6	Need for data transfer,	Expected

Fall 2021	University (SLU)		Staff, GPN CyberTeam		Virtual	network upgrades	Summer 2022
Virtual - 2020/2021	Univ South Dakota	Priv	Staff from GPN, USD, SDSU, Black Hills State, and others	Req	40 Virtual	Regional compute and storage, CI expertise, software integration	Published October 2021
In Person - March 2022	LEARN/McLennan Community College	Priv	Staff from LEARN & McLennan Community College	Req	10 In-Person	Storage and IT staff assistance with R&E workflows	Expected Summer 2022
In Person April 2022	LEARN/Midland College	Priv	Staff from LEARN & Midland College	Req	10 In-Person	TBD	Expected Summer 2022
In Person April 2022	San Diego State University	Priv	Staff from SDSU	Req	10 In-Person & Virtual	TBD	Expected Summer 2022
Virtual - Summer 2021	NOAA Marine Mammal Acoustics (NOAA NMFS)	Priv	Staff from NOAA	Req	10 Virtual	Storage, cloud peering, lack of IT support for CI engineering	Published August 2021
Virtual Spring 2021	University of Central Florida (UCF)	Priv	Staff from UCF	Req	20 Virtual	Need for new CI designs on campus, lack of secure computing, storage	Published August 2021
N/A	Oregon State University	Priv	N/A	Offered	N/A	N/A	Abandoned 2021
N/A	Texas A&M University San Antonio (TAMUSA)	Priv	N/A	Offered	N/A	N/A	Abandoned 2021

5.F Year 5 Plans for Deep Dives

EPOC continues discussion with several parties about Deep Dive activities in Year 5:

- **San Diego State University:** EPOC is working to complete a hybrid Deep Dive with San Diego State University. Virtual components started in early 2022 and will complete with an in-person meeting in April of 2022.
- **LEARN**
 - **Midland College:** EPOC is planning an in-person review with Midland in April 2022. A report will follow this.
 - **St. Mary's College:** EPOC is discussing scheduling this event in late 2022 to coincide with SC22 in Dallas Texas.
 - **South Texas College:** EPOC has indicated to LEARN that this event should be scheduled in early 2023.
- **OARnet:** OARnet has several additional Deep Dives to support a CC* grant, but has struggled to find staff to assist in running them. EPOC is having an ongoing conversation on scheduling.

- **Mississippi State University:** Staff at MSU are interested in a Deep Dive but have not been able to find a time to schedule this event in 2022.
- **Yale University:** Staff at Yale are interested in a Deep Dive but have not been able to find a time to schedule this event in 2022.
- **National Institute of Standards and Technology (NIST):** NIST is interested in hosting a virtual Deep Dive in late 2022 and has been talking to EPOC regarding the specifics.

6. NetSage Deployments and Analysis

Understanding application performance and network measurement are two sides to a single coin - one doesn't make sense without the other. The EPOC project uses the NetSage tool (<http://www.netsage.global>) to collect and evaluate common network measurement data. The initial NetSage software was developed and deployed on the NSF-funded international networks. It was meant to work with sparse, international circuits, and for end users primarily consisting of circuit owners and operators. EPOC has expanded the use of this software to work with more densely defined networks and to support additional analysis and visualizations, and data for all of the NetSage deployments are now available online at <http://all.netsage.global>. More information about NetSage and EPOC is online at <https://epoc.global/wp-content/uploads/Network-Analysis.pdf>. Jennifer Schopf and Dan Doyle jointly lead this activity.

6.A NetSage Development and Presentations

Different components of NetSage can be deployed in different ways, depending on the requirements of the customer. During Year 4, the NetSage development team released versions 1.8.0, 1.8.1, 1.8.3, 1.8.4, 1.9.0, and 1.10.0. Each of these versions were deployed on all of the EPOC partner deployments. All of these releases included numerous bug fixes and minor adjustments. A summary of key changes in these releases is as follows:

- NetSage 1.8.0 fixed a subtle but critical error that would sometimes result in flows being incorrectly shown in an individual flows context.
- NetSage 1.8.1 was a minor fix with some internal cross-linking as a result of changes in 1.8.0.
- NetSage 1.8.3 and 1.8.4 introduced several fixes in different dashboards related to queries designed to show a single flow which were sometimes failing. This was most noticeable for organizations with a very small number of flows.
- NetSage 1.9.0 focused on small but important user experience upgrades, such as more friendly names for source and destination scope selections on the Advanced Flow Analysis Dashboard and consistent ordering of menus across all Dashboards. Additional information designed to help better visualize groupings of flows was added to several Dashboards as well.
- NetSage 1.10.0 included an overhaul of the mapping technology to better support more topologies such as redundant links, along with improved legends and new styling.

Several presentations were made during Year 5. Schopf discussed NetSage as part of the WOMBIR meeting [8]. Schopf also made a presentation [85] to the Sun Corridor Network

Engineering meeting [16] on how they can make better use of their NetSage deployment, and Southworth presented [121] at the SoX engineering meeting [55] with a similar talk. Southworth gave a lightning talk [114] at the Internet2 TechExtra Conference December [46] on NetSage and another presentation [125] at the GÉANT 2nd Performance Management Workshop [59]. Schopf gave an in depth talk to the Routing Working Group on how to use NetSage to find possibly ineffective routes [85], and Meade gave a talk about how the RWG uses NetSage as part of the CI Engineering Lunch and Learn series [100].

A new NetSage Portal was developed for the LHCOPN-LHCONE meeting in October. This portal was designed to showcase existing traffic from LHC, CMS, and ATLAS sites in a project focused manner in order to promote the use of NetSage as a monitoring tool for the scheduled High Luminosity LHC Data Challenge runs. The new portal can be found here: <https://lhc.netsage.global>. This was discussed at several LHC meetings [11, 35, 63], with Southworth giving LHC NetSage presentations [74, 101]

6.B Current Deployments

The status of the current deployments for NetSage network-related dashboards for the EPOC partners includes:

- **Front Range GigaPop (FRGP):** The FRGP flow data deployment of NetSage was made public in December 2019 at <https://frgp.Netsage.global/>. The SNMP data additions were completed October 2021, and the corresponding dashboards are now available.
- **Great Plains Network (GPN):** The NetSage SNMP and flow dashboard for the GPN associated circuits (<http://gpn.Netsage.global>) was initially deployed in October 2018 for SNMP data and in May 2020 deployed flow data collectors as well. In 2021, we worked with GPN to further identify more of their smaller institutions by integrating Shared Whois Project (SWIP) data into the netflow ingest pipeline.
- **iLight/Indiana GigaPop:** Flow data collection for the five Indiana GigaPop routers began in mid-April 2019, and continues to be publicly available at <http://ilight.Netsage.global>.
- **KINBER:** Collection of flow data for the PennREN network began at the end of October 2019 and remains publicly accessible at <https://kinber.Netsage.global/>. As KINBER shifts ownership of its assets to First Light their NetSage setup may need to be adapted. In November 2021, one of the PennREN routers was deconfigured to send data. We have been reaching out to them to coordinate longer term plans as part of the ownership change and will update as the situation evolves.
- **LEARN:** At their 2019 All Hands Meeting, LEARN staff expressed an interest in moving forward to deploy NetSage for the state of Texas network. During Quarter 3, an initial deployment of netflow collection from 11 routers was established within the LEARN network and a corresponding NetSage instance is now available at <https://learn.netsage.global/>. We continue to work with the engineers there to refine and expand this deployment and anticipate new routers will be available.
- **OARnet:** We met with OARnet's new CEO in January 2021. The CEO stated that a NetSage deployment for the Ohio R&E network could not go forward until they had done

a network redesign (currently in the planning stages) to split R&E traffic off from commercial traffic.

- **PIREN:** When PIREN became a member of EPOC, we began to support their deployment of NetSage, previously set up by the IRNC NetSage project. They share a portal with the APOnet collaboration, which is available at <https://aponet.netsage.global>
- **PNWGP/Pacific Wave:** The PNWGP/Pacific Wave deployment of NetSage, which had been supported by the NSF IRNC program, is now being supported by EPOC. This data is available at <https://pacwave.netsage.global/>.
- **SoX:** The SoX NetSage deployment for flow data started receiving data at the very end of June, and was made public in July 2020 at <https://sox.netsage.global/>.
- **Sun Corridor Network:** An initial NetSage deployment for Sun Corridor was made public in March 2021 at <https://suncorridor.netsage.global/>. They expressed interest in using SWIP and adding SNMP data in the future, but this will likely wait until their network re-configuration in early 2022. Due to this network reconfiguration, all data collection efforts have been temporarily paused as of December 2021.
- **TACC:** TACC Flow data has been available since July 2020 and is accessible at <https://tacc.netsage.global/>.

The Archive site deployment was funded by the NSF IRNC NetSage project but is also being used by the various EPOC partners. NetSage uses a software package called Tstat to collect flow data as well as retransmits from the archives. The deployments include:

- **TACC/LEARN:** The TACC deployment remains active, though sometimes requires working with them to restart it based on changes in their environment. No major changes have been made on the EPOC side and this work is running in a stable state.
- **University of Hawai'i Astronomy:** A temporary installation of Tstat for the Astronomy archives was replaced with a permanent solution early in Year 2. This work is running in a stable state.
- **NCAR/FRGP:** A Tstat archive was sent to the lab at NCAR's site in Wyoming in early 2019 and was up and running in July. It has since been running and providing data stably. This work is running in a stable state.
- **National Energy Research Scientific Computing Center (NERSC):** NERSC was the first deployment for IRNC NetSage and the Tstat software. This archive is widely used internationally and domestically for energy science related data sets. This work continues to run in a stable state, periodically communicating with them to coordinate upgrades or resolving issues.

6.C Metrics

Table 3: Metrics for NetSage activities for Year 4

Where Regional	Data	Date Live	# Monitored Devices	# Large Flows	# Unique Src Orgs	# Unique Dest Orgs
FRGP	Flow	1/20	1 router	422,764,286	8,440	12,147
GPN	SNMP, Flow	10/18	2 routers	1,364,921,712	9,945	12,810
iLight	Flow	4/19	5 routers	265,685,463	6,766	15,066
KINBER	Flow	11/19	2 routers	72,804,333	3750	4,495
LEARN	SNMP, Flow	10/21	24 routers	324,416,701	8,578	8,950
PIREN	SNMP, Flow	4/19	2 routers, 1 head node	18,220,521	902	2,294
PNWGP	SNMP, Flow	2/18	7 routers	245,027,984	4,045	7,598
SoX	Flow	7/20	3 routers	103,832,268	4,006	3,575
SCN (temporarily on hold)	Flow	2/21	4 routers	567,935,487	8,445	14,600
TACC (LEARN)	Flow	7/20	1 router, 4 head nodes	57,804,558	2,841	4,957
TACC (LEARN)	Tstat	1/19	6 head nodes	3,986,875	151	155
UHawaii Astro	Tstat	5/19	1 DTN	5,907,908	263	1,195
NCAR (FRGP)	Tstat	7/19	1 DTN	28,083,778	549	2,306
NERSC	Tstat	3/18	10 head nodes	6,968,019	158	175

6.D Network Performance Detection

In Year 4, we continued our efforts to use NetSage to actively investigate performance issues for data transfers across all science domains. This is an extension of the original project milestone of using NetSage to detect or analyze network “disturbances”. We are using NetSage to look at patterns in data movement, volume, and rate to identify and analyze flows with suspected poor data transfer performance or unexpected routing. If disturbances are found we then engage with partner networks and end users to see if we can work together to improve performance. In Year 4, this included Roadside Assistance Cases 160, 164, 165, and 204. We will continue this work in Year 5, as NetSage deployments continue to increase and more flow data becomes available for analysis.

6.E Year 5 Plans

During Year 5, we anticipate that the NetSage Netflow Pipeline version 2.0 will be released. This version includes a significant change to one of key underlying technologies used to collect flows from network devices. As with all pipeline releases, significant work has gone into containerizing the code for third party deployments as well as supporting documentation.

Many of the EPOC partners today run this containerized version of the NetSage Netflow Ingest Pipeline. We anticipate working with all of these partners to upgrade and verify their installations of the new codebase. Given the breadth of network vendors and export configurations that the EPOC project supports across all of its partners, this may also require some fine tuning of the deployment recommendations, updating of documentation, and bug fixes as we go.

In addition to this new pipeline rollout, we will also need to stay on top of regular care and feeding. We plan to keep all of the NetSage dashboards up to date with all future releases, including the release to coincide with the pipeline upgrades. Additionally, we will continue to monitor and work with partners as their networks evolve or to resolve any issues with regular data ingestion.

7. Data Mobility Exhibition/Baseline Performance Testing

The Data Mobility Exhibition (DME) provides institutions with a well-defined, neutrally operated, testing platform of distributed test sites that reports a measurable baseline of data transfer performance. Over 100 institutions have been awarded NSF CC* infrastructure grants to develop Science DMZs or upgrade network bandwidth capacity, however, they have not previously had a way to see the external effects of these deployments. EPOC is currently recommending that institutions aim to be able to transfer **one TeraByte (TB) of data in an hour**, which is equivalent to 2.22 Gb/s average network throughput on a 10G connected host. Those institutions that cannot achieve this are candidates to work with EPOC to determine bottlenecks in their path. The current DME testing framework consists of five test sites hosted at R&E connected institutions, as well as two sites hosted by commercial cloud providers. All seven test sites have equivalent datasets for transfers that range in number from a single file to over 100,000 files, and with file sizes that vary from 100MB to 5TB. Ken Miller is the lead for this area. More information is available at: <https://fasterdata.es.net/performance-testing/2019-2020-data-mobility-workshop-and-exhibition/2019-2020-data-mobility-exhibition/>.

7.A Technical Updates

A number of changes were implemented to the DME testing framework during the last year. These changes were due to some of the R&E connected test points being decommissioned, some new ones being brought online, and others needing to be re-configured after experiencing partial failures:

- Three of the test points were re-configured, after it was found that their test files were being removed from the platform.
- A new test point was added by ESnet, and is located in Washington DC.
- An older test point was decommissioned by ESnet, and will be re-added when the hardware is replaced.

Separate from the work with the hardware platform, ESnet and Globus have developed a set of DME automation scripts to assist with testing, and are available online:

<https://github.com/vasv/dme-utils>. These scripts provide a way to do automated testing via a command line interface instead of running manual tests through the Globus GUI.

EPOC staff continue to work with members of the Globus team in testing and developing a Docker container that is able to run the Globus Connect Server package. This effort is orthogonal to the goals of DME, but will assist in the rapid deployment of Globus endpoints to support research communities. Docker containers are smaller, and facilitate deploying Globus resources closer to the process of science at a number of NSF and DOE facilities. This work dovetails with the ongoing effort by EPOC staff to develop and deploy Docker containers that are capable of running the Globus Modern Research Data Portal (MRDP). In anticipation of the 2022 Globus World conference, Globus developers are updating the MRDP code base to use a number of new tools, and EPOC will continue to work with a number of our users to deploy this as a way to share research data.

7.B Presentations and Work with End Users

During Year 4, we gave a broad set of presentations that included details of the Data Mobility Exhibition and similar topics of data transfer performance, DTN tuning, a Data Transfer Scorecard for baselining, and various campus data architectures [67, 89, 92, 95, 99, 110, 115, 126]. The goal of the presentations, engagement, and assistance has been to increase the information around moving a Terabyte in an hour as a campus baseline. The Data Transfer Scorecard (<https://fasterdata.es.net/DTN/data-transfer-scorecard/>) was developed to provide a reference for data transfer rates based on different scales depending on the audience EPOC may be engaging with during a consultation. Researchers tend to speak about data transfer rate in data volume over time. We created a row in Terabytes per hour. When engaging network architects, engineers, and administrators, this group speaks about data transfer rates in bits per second. We created a row in Gigabits per second. Another row was created for systems or storage administrators who speak about data transfer rates in bytes per second.

During Year 4, 5,486 tests were performed using the DME test sites. These included tests by member institutions of our Regional Networking Partners, including KINBER, GPN, and SoX. EPOC will be working with sites that request assistance to understand and improve data architecture, tuning, and usability of the resources for science use cases.

The University of Alabama Birmingham (UAB), a member of the EPOC Regional Partner SoX, produced ~8800 tests over the past year and a half, and was contacted to gather more details

about their automated testing, as part of Consultation Case #214. UAB has a research computing internship experience that ran during Fall 2020 and Spring 2021 with a focus on data science. A student from this program was using the DME testing to try to understand the effect that a firewall would have on data transfer performance by running tests from two sites, one on the UAB Science DMZ DTN outside of the firewall and one, a Campus DTN, behind the firewall. The large number of tests brought the work to our attention, and we continue to work with UAB staff to evaluate this approach and the results.

EPOC staff are responsive to requests from DME participants that want to understand ways they can improve their outcomes. These requests can come in the form of consultations on data architectures, tuning, and usability of the resources for science use cases. Overall, the DME activity has resulted in institutions understanding their data transfer behaviors from a higher level view. A recent Consultation Case with EPOC focused on assisting a campus identify research groups on campus that may be involved in data transfer activity to other locations. Using monitoring tools available on campus (e.g., sFlow or Netflow/IPFIX), or tools like NetSage that offer a similar aggregated view on a regional or national scale, it was possible to drill down into campus behavior to understand more about the sources, destinations, and tools being used. In order to support additional understanding of the use of the DME test points, we have added their data to the NetSage Science Registry, and using that tool we can see which institutions are using the DME and their performance, as shown in Figure 1.

7.C Year 5 Plans

In Year Five, EPOC will continue to use the DME to test campus and cloud data architectures for effective transfer rates. The Data Transfer Scorecard will be referenced against DME test data for performance and engage campus when poor performance is logged in Globus or discovered with NetSage.

EPOC staff will continue to work with Globus, and will share the containerized version of the Globus Connect Server, to allow for rapid deployments and a greater participation in DME. Additionally, Globus is developing a new Django container version of the Modern Research Data Portal with input from EPOC staff. This new data portal software can be deployed with the core DME datasets, and will serve as an example of a custom data portal for a scientific project or research collaboration. The end goal of this work will be to facilitate the sharing of scientific data sets with collaborators, while still affording access to core functionality (e.g., API access, Identify Access Management (IAM), and other Globus services). This new data portal can provide a logical separation of data portal services from the physical DTN allowing new security policies to be placed around each component separately.




Figure 1: An example of the top DME test point data transfers, as shown in NetSage at https://all.netsage.global/grafana/d/ie7TeamGz/flow-data-for-projects?orgId=2&var-project_type=Data%20Mobility%20Exhibition&var-sensors=All&var-is_net_test=yes&from=1625112000000&to=1633060799000.

8. Managed Services (aka “In a Box”)

EPOC is developing a set of service definitions for common R&E infrastructure components that could be run by a third party as a managed service. The goal of these definitions is to provide guidance for our Regional Networking Partners to implement, maintain, and operate (potentially for a fee) the service as a benefit for downstream connectors. In doing so, the costs associated with design, specification, and installation could be ameliorated for a larger population than would otherwise have access to this technology due to the burdens of entry which may include not having knowledgeable staff or enough compelling use cases to invest time and money. More information about the Managed Service activity is available online at <https://epoc.global/wp-content/uploads/Managed-Services.pdf>. Jason Zurawski is the lead for this area.

The typical design and implementation involves an EPOC regional partner expressing interest in working on managed services to be deployed/managed centrally or targeted directly at a

member school. However, due to pandemic-related access restrictions and the resulting shifting priorities at our partner institutions, these efforts were stalled in mid 2020 and for the foreseeable future.

The only Managed Service in active development at this time is the Modern Research Data Portal (MRDP), as detailed at <http://es.net/science-engagement/technical-and-consulting-services/modern-research-data-portals/>. The MDRP design pattern makes use of the Science DMZ model and DTNs to scale up the data transfer functionality of a data portal. EPOC is experimenting with the concept of a self-contained Data Portal to assist scientific data sharing needs. The goal is to create an easy to install set of software that can be run on campus or regional hardware and expose a set of scientific data.

The current Pilot Portal is based on MRDP and uses Docker and supports a front end with federated authentication and a custom pointer file to any existing site DTN. This setup installs a custom project or site front end with an API call to existing and collaborative Globus end points. Globus has now containerized the Globus Connect Server (<https://github.com/globus/globus-connect-server-deploy>) and included a number of other automation options that will make the operation of this service more straightforward. EPOC continues to engage with Globus, and several partners, on deployment strategies.

In Year 4, we began to work with Baylor University again to install and evaluate a python-based django framework instance of the MRDP. We had previously worked with the same research team from Baylor, that had been identified in the 2020 Baylor Deep Dive as needing a resource like the MRDP. However, they found the previous version of the tool difficult to install and set up. We are hoping that the Baylor team will be able to give us feedback on deployment and documentation as we work together to provide this service.

9. Training

EPOC is continuing the successful training that ESnet and IU led as part of the Operating Innovative Networks (OIN) program. This includes training for network engineers to be coordinated with existing cyberinfrastructure support teams. While training programs like OIN emphasized the design and deployment of technology, we have pivoted to train staff on the use of these tools and the improvement of scientific use of networks through them. In addition to training on tools such as perfSONAR, we offer training for network engineers on interacting with their researchers through teaching them how to perform Deep Dive activities. All EPOC training materials are available online, including lecture materials, exercises, and recorded sessions, when possible.

9.A Collaboration with University of South Carolina

EPOC continues to work with members of the University of South Carolina Cyber Training team, funded via NSF #1829698. EPOC staff are assisting with a set of workshops on high

performance networking technologies. Zurawski is also serving on the Advisory Committee for this project.

The audience for the training includes IT educators, IT professionals, CI engineers, high performance computing specialists, research systems administrators, and security professionals. Topics include science DMZs, TCP, BGP, perfSONAR, and Zeek. One facet that makes this training unique is the use of a virtual laboratory environment that allows for easy set-up, scaling, and quick creation of custom training scenarios using open source tools. Students engage in hands-on training exercises that simulate real world networking leading to better understanding of complex topics. There are over 50 virtual labs associated with the live lectures, and the labs are meant to be completed by attendees at their own pace after a workshop ends.

EPOC staff have worked with USC at the following events during Year 4:

- EPOC staff participated in the virtual USC Cyber Training Workshop [4] sponsored by Southern Crossroads / Southern Light Rail (SoX / SLR), the Engagement and Performance Operations Center (EPOC), the Research Computing (RC) Group at University of South Carolina, Clemson University, Oklahoma University, Alabama A&M University, and the University of South Alabama on April 8, 2021, and April 15, 2021 [64, 65, 66, 67]. The workshop had over 100 registered attendees for the 2-day meeting. Audience members included engineers, higher education staff, and college educators.
- EPOC staff participated in a virtual workshop [54] in February 2022 co-sponsored by GPN. Topics were advanced, and focused on P4, a programming paradigm for software-defined networks. Meade and Addleman led a session out routing and performance [120].

EPOC staff plan to continue to work with USC plan events during Year 5, including:

- EPOC staff will participate in a virtual workshop in April 2022 co-sponsored by NYSERNet. Topics include Science-DMZ, Data Mobility, and Monitoring.
- EPOC staff are discussing possibilities for an event with CENIC, co-located with their annual meeting, in September 2022.
- EPOC staff are discussing possibilities for an event with the NSF and the Quilt, co-located with their annual meeting, in September 2022.
- EPOC staff are discussing possibilities for an event with Internet2, co-located with their TechX event, in December 2022.

9.B Border Gateway Protocol (BGP) Training

EPOC has received several requests for Roadside Assistance or Consultations that have involved correcting the configuration and management of Border Gateway Protocol (BGP), particularly when an institution is balancing traffic between R&E networks and commodity networks. The problems our end users are experiencing are not related to setting up an initial instance of BGP, but in making the correct adaptations to the BGP tables as capacity is added in order for data flows to still be routed effectively. While EPOC will continue to work with sites on a case-by-case basis to explain and fix this type of configuration and operational issue, it

was determined that a more focused effort was needed to support the development and dissemination of educational materials to instruct and explain BGP adaptations. Our work with BGP is complemented by the GNA-G / APAN Routing Working Group, discussed in Section 3.F.

9.C PerfSONAR Training

In addition to the Quarter 1 training that was part of the USC Cyber Training Workshop, discussed in Section 9.A, EPOC staff lead perfSONAR training sessions at the GPN annual meeting [13, 80] and the LEARN Technical Advisory Group [14, 81]. These sessions focused on using perfSONAR to detect and troubleshoot the types of complex network problems that arise as R&E transfers cross multiple network domains, as well as the importance of maintaining healthy, available perfSONAR resources for others in the R&E community to use. While no presentations or trainings were scheduled in Quarters 2-4, Southworth has been invited to lead a half-day workshop at the GPN Annual Meeting in June 2022. This workshop will include the use of NetSage to find end-to-end transfer performance issues, and advanced debugging of these issues using perfSONAR.

EPOC staff use perfSONAR extensively as part of the Roadside Assistance and Consulting project, as detailed in Section 4. As such, outreach efforts that increase perfSONAR presence and knowledge are valuable to both EPOC and the larger R&E community as a whole. We also participated in the Second European perfSONAR User Meeting in April [6, 68], the perfSONAR All Hands Meeting in July [18], and the TechEXtra21 perfSONAR Day [43].

10. Data Privacy and Security

Basic security measures are being maintained and there were no security incidents to report for Year 4.

Continuing training for Addleman, the Security Officer for the team, took place in the form of attending the Large Scale Facility meeting [1], the CyberSecurity as Big Data Science workshop [5], the Internet2 Online session on Identifying Cyberinfrastructure Gaps [9], the Internet2 Network Security at the Border Workshop [24], the Trusted CI 2021 Cybersecurity Summit held virtually in October [37], and the IU Research Security Operations Center (ResearchSOC) Cybersecurity Engagement in a Research Environment Workshop in December [47].

No PII is shared in the Roadside Assistance or Consultation summaries or reports, which are made public. There may be PII in other documents in a Roadside Assistance Case Folder, for example IP addresses, but this information is locked down and access is controlled and only shared with specific staff working on a particular issue.

In addition, NetSage does not collect PII. The IRNC NetSage privacy documentation is available online at <https://epoc.global/wp-content/uploads/Data-Privacy-Policy.pdf>.

11. Travel and Virtual Meeting Participation

EPOC staff participated in various meetings to support ongoing deployment, collaboration, and training. Starting in February 2020, activities that involved travel were severely impacted by COVID-19. The EPOC activities shifted from in-person to remote/virtual interactions.

During the Year 4, meeting participation by the team listed here and referred to throughout the report with the reference number listed:

1. Addleman attended the Large Facilities Workshop, April 7, 2021, <https://www.largefacilitiesworkshop.com/apr21workshop/>. Addleman participated in discussions on the challenges faced during the COVID-19 pandemic and how public outreach helped continue missions virtually.
2. Schopf attended the Spring CASC 2021 Membership Meeting, April 7-9, 2021, <https://casc.org/event/casc-spring-2021-membership-meeting/>. Schopf participated as a role member of the Executive Committee.
3. Addleman attended the FABRIC Experimenters Workshop, April 8-9, 2021, <https://whatisfabric.net/events/fabric-experimenters-workshop-2021>. Addleman participated in discussions on FABRIC installation status, talks by experimenters planning to use FABRIC, and a demonstration showing how experiments can use FABRIC.
4. Zurawski, Miller, Robb, and Southworth presented the USC Cyber Training Workshop April 8, 2021 & April 15, 2021, http://ce.sc.edu/cyberinfra/nsf_cc_workshop.html. Zurawski, Miller, Robb, and Southworth presented on data architecture, perfSonar, DTNs, and data mobility.
5. Addleman attended the Cybersecurity as Big Data Science Interactive Workshop, April 12, 2021, <https://nebigdatahub.org/cybersecurity-as-big-data-science-workshop/>. Addleman participated in discussions of both common and unique data science challenges for cybersecurity including privacy, collection, storage, and sharing of data.
6. Southworth presented at the Second European perfSONAR User Workshop, April 14 2021, <https://wiki.geant.org/display/perfSONAR/2nd+European+perfSONAR+User+Workshop>. Southworth gave an overview of two success stories which demonstrated the day-to-day usefulness of perfSONAR as a network troubleshooting tool for complex networking problems.
7. Moynihan attended the VRO-NET meeting, April 15-16, 2021. Moynihan participated in discussions on trans-Atlantic networking support for the VRO, contributed to writing documentation on international connectivity, and discussed how EPOC could support the end-to-end data transfers domestically.
8. Schopf attended the NSF Workshop on Overcoming Measurement Barriers to Internet Research (WOMBIR 2021), April 15-16, 2021, <https://www.caida.org/workshops/wombir/2104/>. Schopf participated in discussions around internet measurements and use cases.
9. Addleman attended the Internet2 Online session on Identifying Cyberinfrastructure Gaps, April 20, 2021, <https://internet2.edu/i2-online-recommendations-identifying->

- [cyberinfrastructure-gaps/](#). Addleman participated in a discussion about the barriers to research computing for campuses.
10. Zurawski attended the 2021 KanREN Annual Meeting, April 21-22, 2021, <https://www.kanren.net/annual-meeting/>. Zurawski participated in discussions about research and institutional network needs.
 11. Moynihan and Southworth attended the High Luminosity Large Hadron COLLider (HL-LHC) Data Challenge Monitoring Workshop, April 27, 2021, <https://indico.cern.ch/event/1027287/>. Moynihan and Southworth participated in several discussions about monitoring solutions for the upcoming LHC Data Challenge later this year. Southworth presented on how NetSage could be used to provide real-time and post-event analysis.
 12. Zurawski, Miller, and Robb attended the GlobusWorld 2021 event, May 12-14, 2021, <https://www.globusworld.org/conf/>. Robb presented on recovery work for the Arecibo Observatory.
 13. Schopf, Addleman and Southworth attended the GPN 2021 Virtual Annual Meeting, June 2-4, 2021, <https://www.greatplains.net/gpn2021-virtual-annual-meeting/>. Schopf and Addleman lead a BOF on how EPOC supports the NSF CC* program. Southworth held a perfSONAR training workshop.
 14. Southworth attended the LEARN Technical Advisory Group (TAG) Meeting, June 8-10, 2021. Southworth held perfSONAR training workshop sessions during the meeting.
 15. Schopf, Addleman and Meade attended the GNA-G Community VCs Meeting, June 14-15, 2021, <https://www.gna-g.net/attend-a-meeting/gna-g-community-vcs-2021q2/>. Schopf presented on the status of the Routing Working Group.
 16. Schopf presented at the Sun Corridor Network Engineering Meeting, June 14, 2021. She gave a presentation on NetSage and how it could be used by Sun Corridor and its members.
 17. Moynihan and Schopf attended the TNC 2021 event, June 21-25, 2021, <https://tnc21.geant.org>. Moynihan participated in several update sessions including on AfricaConnect3, NORDUnet, and GEANT. Schopf led a BOF session on the new Routing Working Group.
 18. Moynihan attended the Pacific Research Platform (PRP) Capstone event June 22, 2021, https://drive.google.com/file/d/18njsRQ7IFjMNq1_a_IlePnHfrvjLWEZG/view. Moynihan participated in talks on the history, accomplishments, and the future plans of the PRP.
 19. Southworth attended the perfSONAR developer All Hands Meeting on July 13-15, 2021. Southworth participated in discussions on a variety of topics related to current release status, future releases, and training efforts.
 20. Jennifer Schopf attended the PEARC'21 Conference on July 19-22, 2021, <https://pearc.acm.org/pearc21/>. Schopf participated in several sessions, including several related to CaRCC.
 21. Moynihan attended the PRAGMA Reconnect Webinar on July 22, 2021. Moynihan presented on the EPOC end-to-end performance work and participated in discussions to find new areas of collaboration for the PRAGMA community.
 22. Schopf attended the APAN 52 Conference, August 2-6, 2021, <https://apan52.apan.net/wp/>. Schopf gave an Invited Keynote about the APOnet

- collaboration, presented at the Network Engineering Working Group, and ran a BoF session for the Routing Working Group.
23. Zurawski attended and presented virtually at the NOAA N-Wave Joint Engineering and Technical Interchange (JETI) Conference, August 4-6, 2021, <https://www.noaa.gov/organization/information-technology/n-wave-joint-engineering-and-technical-interchange-jeti>. Zurawski presented on the recent NOAA Deep Dive report.
 24. Addleman attended the Internet2 Community Voices Series: Cost-Effective, Scalable Solution for Network Security at the Border, August 12, 2021, <https://internet2.edu/internet2-community-voices-series-cost-effective-scalable-solution-for-network-security-at-the-border/>. Addleman participated in talks centered on network border security that uses a combination of software tools and network hardware to automate security filters.
 25. Miller attended and presented at the Globus Webinar on Globus and the Science DMZ on August 12, 2021, <https://www.globus.org/events/globus-and-science-dmz-indispensable-tools-cc-awardees>. Miller presented on the DME.
 26. Robb attended the University of Central Florida IT Webinar, on August 13, 2021, and presented about the Arecibo Observatory Data Movement project.
 27. Schopf attended the Internet2/CarCC Online Supporting Data-Driven Discovery on Campus on August 18, 2021, <https://internet2.edu/i2-online-supporting-data-driven-discovery-on-campus/>. Schopf participated in discussions focused on building community along various tracks that reflect the RCD professional activities and how CarCC is working to increase visibility and participation, particularly among underserved organizations.
 28. Addleman attended and presented at the I-Light Town Hall Meeting on August 31, 2021. Addleman presented an overview of the EPOC project.
 29. EPOC staff attended the NSF CC* PI Workshop on September 14-15 and 21-22, 2021, <https://www.thequilt.net/public-event/2021-nsf-virtual-cc-pi-workshop/>. Addleman and Meade, with Cas D'Angelo from SoX, facilitated a breakout session on changes to technology plans based on technology requirements and equipment availability. Miller and Southworth hosted Drop-In Office Hours for NSF CC EPOC which included presenting a DME overview. Zurawski and Addleman hosted Drop-in office hours for EPOC as well and discussed the Deep Dive process with attendees. Zurawski also led an Overview of EPOC's Deep Dive Research Engagement Process.
 30. Addleman attended the FAB All Hands Meeting on September 16, 2021. Addleman participated in discussions around the current status of connecting the University of Bristol to the FABRIC network in the US. FABRIC staff also gave an update on the current FABRIC network and hardware deployment.
 31. Moynihan attended the Second Global Research Platform (2GRP) Workshop on September 20-21, 2021, <https://grpworkshop2021.theglobalresearchplatform.net/>. Moynihan participated in partner update sessions and sessions on NSF-funded projects FAB and BRIDGES.
 32. Addleman, Zurawski, Meade, and Southworth attended the Quilt Fall Members Meeting on September 29-30, 2021, <https://www.thequilt.net/public-event/the-quilt-2021-virtual-fall-member-meeting/>. They participated in talks focused on the MANRS project, a

review of the NSF CC* PI Workshop, and the future of fiber purchases and inter-exchange point links.

33. Moynihan and Meade attended the I-Light Members Meeting on September 28-29, 2021, <https://ilight.net/members-meeting/agenda/>. They participated in sessions on cybersecurity, research support, and future I-light activities.
34. Miller attended the Fall SoX Participant Meeting, October 7, 2021, <https://www.sox.net/2021-fall-participants-meeting/>. Miller presented on DME, DTNs, Data Architectures, and an EPOC overview.
35. Addleman, Moynihan and Southworth attended the LHCOPN-LHCONE Meeting #47, October 11-12, 2021, <https://indico.cern.ch/event/1022426/>. Southworth presented a proposal for using NetSage to monitor LHC traffic in upcoming data challenges.
36. Zurawski attended the BRICCS Landscape Workshop 2021, October 18, 2021, https://hprc.tamu.edu/briccs/workshop_2021.html. Zurawski presented in the session on network consulting offerings.
37. Addleman attended the 2021 Cybersecurity Summit, October 18-20, 2021, <https://www.trustedci.org/2021-cybersecurity-summit>. Addleman participated in discussions on required security levels of an NSF project, best practices for securing a ScienceDMZ, and security impacts on network architectures.
38. Southworth attended the GlobalNOC Connects Annual Members Meeting, October 19-20, 2021, <https://usergroup.globalnoc.iu.edu/events/globalnoc-connects-2021/>. Southworth participated in discussions about infrastructure updates and plans for future product development by the Global NOC and gave a lightning talk on EPOC.
39. Schopf attended the CASC Fall Meeting, October 19-21, 2021. Schopf participated in her role as Executive Committee Member, and pointed several attendees to EPOC materials when relevant.
40. Dart and Meade attended the CaRCC Systems Facing Track Call, October 21, 2021, <https://carcc.org/event/systems-facing-track-call-october-21st-2021/>. Dart and Meade presented two talks, one on network troubleshooting and a separate EPOC overview.
41. Southworth attended the 2021 Regional Community Meeting of the Midwest Big Data Innovation Hub, October 27-28, 2021, <https://midwestbigdatahub.org/2021-rcm/>. Southworth participated in discussions that focused on ways to improve community and data portability between researchers.
42. Schopf attended the SHARP CCI Meeting, October 29, 2021. Schopf presented on EPOC and support for CI Plans at the meeting which was part of a recently funded CC* project at University of Arkansas.
43. Southworth attended the TechEXtra21: perfSONAR Day on November 3, 2022, <https://internet2.edu/i2-techextra/techextra-schedule/perfsonar-day-2021/>. Southworth participated in discussions about current and future perfSONAR development.
44. Schopf, Addleman, Zurawski, and Meade attended SC'22, November 15-18, 2022, <https://sc21.supercomputing.org/>. Addleman, Zurawski and Meade were part of the SCinet team. Schopf attended virtual sessions. Meade also accepted the HPCwire Readers' Choice Award on behalf of the EPOC team.
45. Zurawski participated in a virtual Deep Dive with South Plains College, jointly with LEARN, November 19, 2021. The engagement provided a number of findings and the

EPOC team has made recommendations to the South Plains and LEARN teams for strategic planning.

46. Schopf, Addleman, Southworth, Meade and Miller attended the Internet2 TechExtra Conference December 1-3, 2021, <https://internet2.edu/techextra21-infrastructure-advanced-networking/>. Addleman presented on the success of the EPOC Roadside Assistance and highlighted interesting cases worked by the team including the Arcibo data transfer to TACC. Meade presented a lightning talk on the Routing Working Group highlighting topics of the meetings, cases over the past 6 months and an overview of how to get involved in the Routing Working Group. Southworth presented a lightning talk on the future of the NetSage project. Miller presented a lightning talk on the DME.
47. Addleman attended the Cybersecurity Engagement in a Research Environment Workshop, December 7-9, 2021, <https://blogs.iu.edu/researchsoc/2021/11/11/workshop-cybersecurity-engagement-in-a-research-environment/>. Addleman, to continue network security training, participated in discussions on the security needs of researchers versus enterprise systems, and how to adapt security policy to researchers and working to meet their needs.
48. Zurawski attended the LEARN Board meeting on December 7, 2021, and answered questions related to EPOC Deep Dives supporting the region.
49. Schopf attended the KNIT Winter '21: A FABRIC Community Workshop, December 8, 2021, <https://whatisfabric.net/events/knit-winter-2021-fabric-community-workshop>. Schopf participated in discussions about the structure of FABRIC and how EPOC might be able to support the project.
50. Zurawski participated in a virtual Deep Dive with Sinclair College, jointly with OARnet, December 9, 2021. The engagement provided a number of findings and the EPOC team has made recommendations to the ASU team for strategic planning
51. Schopf attended the Fall AGU Meeting, December 13-17, 2021, <https://www.agu.org/Fall-Meeting-2021>. Schopf met with geoscience colleagues and left 100 reach outs for EPOC help on posters over the four-day meeting.
52. Zurawski and Meade participated in a virtual Arizona State University Deep Dive, Tempe, AZ, January 10-12, 2022. The engagement provided a number of findings and the EPOC team has made recommendations to the ASU team for strategic planning.
53. Zurawski and Addleman participated in a virtual St. Louis University Deep Dive, January 25, 2022. The engagement is the first step in helping SLU design a data mobility strategy.
54. Meade, Addleman, and Zurawski attended the virtual Programmable Switches Workshop, virtual, February 16-17 and 23, 2022, http://ce.sc.edu/cyberinfra/p4_workshop_feb_2022.html, sponsored by GPN, University of South Carolina, NSF, and EPOC. Zurawski was a coordinator for sessions on fundamentals of P4, applications of P4 switches, and hands-on sessions with P4 switches. Addleman and Meade also presented on network performance.
55. Southworth participated in a virtual SoX Engineering meeting, February 24, 2022. He presented a tutorial on NetSage and EPOC roadside assistance.

56. Schopf attended (virtually) and Zurawski attended (in person) the NSF Cyberinfrastructure for Major Facilities Workshop, Redondo Beach, CA, March 2-3, 2022, https://whova.com/web/hecmw_202202/. Zurawski gave a lightning talk on EPOC.
57. Zurawski attended the NOAA Virtual Stakeholders Meeting, March 1-3, 2022. He gave an overview talk on EPOC.
58. Zurawski and Miller attended a meeting with Pennsylvania State University, March 3, 2022. They gave an EPOC overview for their monthly Research State of the Practice Group.
59. Southworth presented at the GÉANT 2nd Performance Management Workshop, virtual, March 08, 2022. Southworth presented on the NetSage monitoring tool.
60. Zurawski and Miller attended the Quilt Spring meeting, March 8-10, 2022, <https://www.thequilt.net/public-event/2022-quilt-virtual-winter-member-meeting/>. Miller gave a DME talk.
61. Zurawski and Meade participated in a McLennan Deep Dive, Waco, TX, March 14, 2022. The engagement provided a number of findings and the EPOC team has made recommendations to the team for strategic planning.
62. Meade and Moynihan attended the GNA-G Community VCs meeting, March 30, 2022, and gave a brief update of the GNA-G Routing Working Group. A high level overview of the ESnet 6 project was given.
63. Moynihan and Southworth attended the virtual LHCONE meeting, March 29-30, 2022. Southworth showed the current progress of NetSage as part of a larger conversation involving monitoring LHC traffic during the data challenges. Moynihan engaged in discussions on setting up deliberate back-ups across the trans-Atlantic links for LHCONE traffic.

12. Presentations and Publications

For Year 4, the EPOC presentations, invited posters, and publications are listed here and referred to throughout the report with the reference number listed:

64. Zurawski, J., “Designing, Building & Maintaining a Data Architecture”, Invited Talk, USC Cyber Training Workshop, April 8, 2021, Slides available online at https://drive.google.com/file/d/1i6fSwsJx8JyD5JtgT_8pU3Z_ZMC2tr_J/view?usp=sharing
65. Southworth, D., “perfSONAR Introduction & Motivation”, Invited Talk, USC Cyber Training Workshop, April 8, 2021, Slides available online at https://drive.google.com/file/d/1jX0JAzDbJswDhdW_GY_aEqC4XJhLReP4/view?usp=sharing
66. Robb, G., “Move That Data! DTN Elements”, Invited Talk, USC Cyber Training Workshop, April 8, 2021, Slides available online at <https://drive.google.com/file/d/1IWfuPXnipY9aQmne6kYedIP4d4Rg3EJ0/view?usp=sharing>
67. Miller, K., “Data Mobility Exhibition – DME”, Invited Talk, USC Cyber Training Workshop, April 8, 2021, Slides available online at

- <https://drive.google.com/file/d/1SteTKEMqe4DzkRV9zvdIwKsdEnAFs1bw/view?usp=sharing>
68. Southworth, D., "In the Wild: Real-world Troubleshooting with perfSONAR", Invited Talk, Second European perfSONAR User Workshop, April 14, 2021. Slides available online at:
<https://wiki.geant.org/display/perfSONAR/2nd+European+perfSONAR+User+Workshop?preview=/138747020/255557730/pS%20in%20the%20wild%20-%20IU-DS.pdf>
 69. Moynihan, E., "Kyungpook National University Data Transfer Performance to CERN", EPOC White Paper, April 16, 2021, Updated in November 2021, Available at:
<https://epoc.global/wp-content/uploads/Kyungpook-National-University-to-CERN.pdf>
 70. Moynihan, E., "Japan to Korea Route Efficiency ", EPOC White Paper, April 20, 2021, Updated November 2021, Available at: https://epoc.global/wp-content/uploads/NII_Korea-Routing-Issue.pdf
 71. EPOC Team, "Continuing Arecibo's Legacy", Press release, April 21, 2021, Available at: <https://www.tacc.utexas.edu/-/continuing-arecibo-s-legacy>
 72. EPOC Team, "Continuing Arecibo's Legacy", Press release, April 21, 2021, Available at: <https://itnews.iu.edu/articles/2021/Continuing-Arecibos-legacy-.php>
 73. EPOC Team, "Arecibo Data Recovery: Behind the Scenes with Jason Zurawski", ESnet Blog post, April 21, 2021, Available at: <https://lightbytes.es.net/2021/04/21/arecibo-data-recovery-behind-the-scenes-with-jason-zurawski/>
 74. Southworth, D., "NetSage- A Tool to Understand Data Transfers", Invited Talk, HL-LHC Data Challenge Monitoring workshop, April 27, 2021. Slides available online at: <https://indico.cern.ch/event/1027287/contributions/4335173/attachments/2234274/3786518/LHC%20NetSage.pdf>
 75. Meade, B., "National Center for Atmospheric Research (NCAR) Multicast Performance", EPOC White Paper, May 2, 2022, Available at: <https://epoc.global/wp-content/uploads/National-Center-for-Atmospheric-Research-Performance.pdf>
 76. Addleman, H., Dart E., Zurawski, J., Southworth, D., and Miller, K., "It's all about the MTU's", EPOC White Paper, First available May 13, 2021, updated August 2022, Available at: <https://epoc.global/wp-content/uploads/About-MTUs.pdf>
 77. Robb, G., "The Arecibo Observatory: Disaster, Recovery and What Comes Next", Invited Talk, GlobusWorld 2021, May 13, 2021, Slides available online at: <https://www.slideshare.net/globusonline/globusworld-2021-arecibo-observatory-data-movement>
 78. Zurawski, J., Schopf, J., "EPOC Trials Virtual Deep Dives to Enhance Cyberinfrastructure Workflows", Invited Article, The Quilt Circle, Pg 13, June, 2021. Article available online at: <https://www.thequilt.net/quilt-circle/epoc-trials-virtual-deep-dives-to-enhance-cyberinfrastructure-workflows/>
 79. Schopf, J., and Addleman, H., "EPOC Support of NSF CC*", Invited BOF, GPN 2021 Virtual Annual Meeting, June 2, 2021, Materials available online at: <https://gpn2021.sched.com/event/jC0F/nsf-cc-bof?iframe=no&w=100%&sidebar=no&bg=no>

80. Southworth, D., "perfSONAR Training", GPN 2021 Virtual Annual Meeting, June 2, 2021, Slides available online at: https://drive.google.com/drive/folders/18NVEB_qm_AxGsD-hlladFIOSlr6F5OzBW?usp=sharing
81. Southworth, D., "perfSONAR Training", LEARN Technical Advisory Group (TAG) Meeting, June 8 & 10, 2021, Slides available online at: <https://drive.google.com/drive/folders/1ZSgwlrzwfZ7gS-NBtLrEc7ejRSI98M-y?usp=sharing>
82. Schopf, J., "Routing WG Overview", Invited Talk, GNA-G Community VCs Meeting, June 14, 2021. Slides available online at: https://drive.google.com/file/d/18jY7-HGpbCam9rhUQP01kxpe_O7PkZ8U/view
83. Schopf, J., "NetSage, EPOC, and the AZ Tri-University Network" Invited Talk, Sun Corridor Network Engineers Meeting, June 14, 2021. Slides available online at: https://drive.google.com/file/d/1RF0_r_q6Lw3BdSgLQ7ZoBPEq4fWSDoMb/view?usp=sharing
84. Schopf, J., "Routing WG Overview", Invited BOF Session, TNC21, June 25, 2021. Slides available online at: https://drive.google.com/file/d/1jEvF_Kb04zZPz1ciaGifsLRc6eUqvfgR/view?usp=sharing
85. Schopf, J., "NetSage and Finding Unexpected Routes", Invited Talk, Routing Working Group, July 20, 2021. Slides available online at: <https://docs.google.com/presentation/d/1x07cL8-o38sWHZd-RXSnZLinkj0t02IBXjpdLa3C45o/edit?usp=sharing>
86. Moynihan, E., "IN@IU and TransPAC Update", Invited Talk, PRAGMA Reconnect Webinar, July 22, 2021. Slides available online at: <https://drive.google.com/file/d/1PBLRdMvZYvjYvNK3H39Ws0SqUZt0Mb6c/view?usp=sharing>
87. Schopf, J., "Routing WG Overview", BOF Talk, APAN 52, August 5, 2021. Slides available online at: https://drive.google.com/file/d/1jEvF_Kb04zZPz1ciaGifsLRc6eUqvfgR/view?usp=sharing
88. Zurawski, J., "EPOC Deep Dive: NOAA National Centers for Environmental Information - Fisheries Acoustics Archive", Invited Talk, NOAA NOAA N-Wave Joint Engineering and Technical Interchange (JETI), August 5, 2021. Slides available online at: <https://docs.google.com/presentation/d/1LQrma77cX7fmGklm7r8y88hGkX9oATyk/edit?usp=sharing&oid=113418983892402433394&rtpof=true&sd=true>
89. Miller "Data Movement Exhibition (DME)", Invited Talk, Globus Webinar on Globus and the Science DMZ, August 12, 2021. Slides available online at: https://drive.google.com/file/d/1apW7JPMWKcWU_628BcPE18YQERggMmDE/view?usp=sharing
90. Robb, G., "Arecibo Observatory Data Movement", Invited Talk, University of Central Florida IT, August 13, 2021. Slides available online at: https://drive.google.com/file/d/1MGisEyTEwOz1_fr4MFymEtZmGKqhfGdF/view?usp=sharing
91. Zurawski, J., Addleman, H., Miller, K., Southworth, D., "NOAA National Centers for Environmental Information Fisheries Acoustics Archive Network Deep Dive", Lawrence

- Berkeley National Laboratory, Report #: LBNL-2001417, August 19, 2021. Available online at: <https://escholarship.org/uc/item/0wm478c4>
92. Addleman, H., "Moving Data Faster with EPOC", Invited Talk, I-Light Town Hall Meeting, August 31, 2021. Slides available online at: <https://drive.google.com/file/d/1G-3RPSd6y3XIQBDpJ5MdxtzYXz46n4B/view?usp=sharing>
 93. Zurawski, J., Addleman, H., Miller, K., "University of Central Florida Campus-Wide Deep Dive", Lawrence Berkeley National Laboratory, Report #: LBNL-2001419, August 31, 2021. Available online at: <https://escholarship.org/uc/item/7fh8p0mw>
 94. Southworth, D., and Miller, K., "Easy 10G Data Transfer Nodes", EPOC White Paper Video, September 8, 2021. Available at: <https://youtu.be/3WiraL5triv>
 95. Miller, K., Southworth, D., "Data Mobility Exhibition (DME): A Data Transfer Performance Framework", Invited Talk, NSF CC* PI Workshop, September 15, 2021. Slides available online at: https://drive.google.com/file/d/1joxc_ZP-0B7MMLVRbFHIzI7KE5oYNob/view?usp=sharing
 96. Zurawski, J., "Overview of EPOC "Deep Dive" Research Engagement Process", Invited Talk, NSF CC* PI Meeting, September 21, 2021. Slides available online at: <https://docs.google.com/presentation/d/1OZg6kKFctY5m3W1Oy7tB04j5b29q1ldv/edit?usp=sharing&oid=113418983892402433394&rtpof=true&sd=true>
 97. Addleman, H., Zurawski, J., "EPOC Deep Dives: What We Hear & How to Handle", Invited Breakout, NSF CC* PI Meeting, September 21, 2021. Slides available online at: https://docs.google.com/presentation/d/11WPgr_GeArZ4QEmZlrptZ0NXFHa8hY3n/edit?usp=sharing&oid=113418983892402433394&rtpof=true&sd=true
 98. Addleman, H., Meade, B., D'Angelo, C., "Changes to Technology Plans Based on Technology Requirements and Equipment Availability", Invited BOF, NSF CC* PI Meeting, September 21, 2021.
 99. Miller, K., "Data Mobility Exhibition (DME): A Data Transfer Performance Framework", Invited Talk, Fall SoX Participant Meeting, October 7, 2021. Slides available at: <https://drive.google.com/file/d/15YKakyAF5chwpapuiznn8DZbrktsak2R/view?usp=sharing>
 100. Meade, B., Moynihan, E., "Investigating, troubleshooting, and improving performance with NetSage and the GNA-G Routing WG", Invited Talk, CI Engineering Lunch & Learn Series, October 8, 2021. Video available online at: <https://www.youtube.com/watch?v=hSmls2YeZ50>
 101. Southworth, D., "Understanding LHC data movement with NetSage". Invited Talk, LHCOPN-LHCONE Meeting #47, October 12, 2021. Slides available online at: <https://drive.google.com/file/d/11Xcl3UJmnb7pxVLATlgHilBMTCj7hOs/view?usp=sharing>
 102. Zurawski, J., "Overview: The Engagement and Performance Operations Center", Invited Talk, BRICCS Landscape Workshop 2021, October 18, 2021. Slides available at: <https://drive.google.com/file/d/1QiXeXn4FVZQex16xajg-E60AASIC44FM/view?usp=sharing>
 103. Southworth, D., "EPOC Overview (It's pronounced EPIC)", Invited Lightning Talk, GlobalNOC Connects Annual Members Meeting, October 20, 2021.

104. Dart, E., "Network Troubleshooting: Techniques and Approaches", Invited Talk, CaRCC Systems Facing Track Call, October 21, 2021. Video available online at: https://www.youtube.com/watch?v=Ezbe42w3d30&list=PLV-SC0CHLTwehApeMg_fJ5VbYndETH3yT&index=4
105. Meade, B., "EPOC Overview", Invited Talk, CaRCC Systems Facing Track Call, October 21, 2021. Slides available online at: <https://drive.google.com/file/d/1EDIZ1q-G8kbMCB3DCwN5FrcwvRqWigOt/view?usp=sharing>
106. Addleman, H., "Arecibo Data Preservation Summary", White Paper, October 28, 2021, Available online: <https://epoc.global/wp-content/uploads/Arecibo-Data-Preservation-Summary.pdf>
107. Schopf, J., "EPOC Overview and Review of CI Plans", Invited Talk, SHARP CCI Meeting, October 29, 2021. Slides available online at: https://drive.google.com/file/d/1T1uNhJV1H_ojYUGqEuEiMFMsDVJ6FQc/view?usp=sharing
108. Zurawski, J., Addleman, H., Miller, K., Southworth, D., "South Dakota Region Scientific Deep Dive", Lawrence Berkeley National Laboratory, Report #LBNL-2001427, November 2, 2021. Available online at: <https://escholarship.org/uc/item/3s07q0d8>
109. Zurawski, J., Addleman, H., "Kennesaw State University Scientific Deep Dive", Lawrence Berkeley National Laboratory, Report #LBNL-2001435, November 3, 2021. Available online at: <https://escholarship.org/uc/item/77k892k5>
110. Miller, K., Robb, G., "Data Mobility Exhibition (DME): Measure Twice, Copy Once", Invited Talk, CI Engineering Lunch & Learn Series, November 12, 2021. Video available at: <https://www.youtube.com/watch?v=ZowMQnykxLI>
111. EPOC Team, "Leading Research Institutions Receive HPCwire Readers' Choice Award", Press Release, November 16, 2021, Available online at: <https://itnews.iu.edu/articles/2021/leading-research-institutions-receive-HPC-wire-readers-choice-award.php>
112. Addleman, H., "Roadside Assistance Year in Review", Invited Talk, Internet2 TechExtra Conference December 2, 2021. Slides available online at: <https://drive.google.com/file/d/1Y4aTbuBokYR3fW19zllglo72O33lnWgd/view?usp=sharing>
113. Meade, B., "International Routing Working Group", Invited Talk, Internet2 TechExtra Conference December 2, 2021. Slides available online at: <https://drive.google.com/file/d/1rPPoVWwKOSFKBb5TYll6xwg6gABjCtkH/view?usp=sharing>
114. Southworth, D., "NetSage: Next Steps", Invited Talk, Internet2 TechExtra Conference December 2, 2021. Slides available online at: https://drive.google.com/file/d/1SYGRIXqFSTkbRANVKOxw_1HGeO1GJC_/view?usp=sharing
115. Miller, K., "Data Mobility Exhibition (DME): Measure Twice, Copy Once", Invited Talk, Internet2 TechExtra Conference December 2, 2021. Slides available online at: <https://drive.google.com/file/d/1aR2YcDsnkgmhliC0NjtxgpBe0hF66g5f/view?usp=sharing>

116. Southworth, D., "DTN Build Basics", EPOC White Paper, December 22, 2021. Available online at: <https://epoc.global/wp-content/uploads/DTN-Build-Basics.pdf>
117. EPOC Team, "EPOC provides personalized technology support for science communities ", Quilt Circle Article, January 2022, Available at: <https://www.thequilt.net/the-quilt-circle/>
118. Abhinit, I., Addleman, H., Benninger, K., DuRousseau, D., Krenz, M., Meade, B., "Science DMZ: Secure High Performance Data Transfer", Indiana University Tech Report DOI: 10.5967/f4aj-t870, January 6, 2022. Available at: <https://epoc.global/wp-content/uploads/Science-DMZ-Secure-High-Performance-Data-Transfer.pdf>
119. Trusted CI Team, "Trusted CI, EPOC and University of Arkansas Create Security Resources for Science DMZs", Blog Post, January 7, 2022, Available at: <https://blog.trustedci.org/2022/01/trusted-ci-epoc-and-university-of.html>
120. Meade, B., Addleman, H., "Routing, Buffers, and Network Performance, Oh my", Training Workshop, University of South Carolina Programmable Networks Training, February 23, 2022. Slides available at: <https://drive.google.com/file/d/1EByeFNKdpQlr-GSumjjeak6dS8HbajvV/view?usp=sharing>
121. Southworth, D., "NetSage, EPOC, and SoX", Invited Presentation, SoX Engineering Meeting, February 24, 2022. Slides available at: https://drive.google.com/file/d/1BEHvmB8KITVxEigbmTUYiLVM_vXM4256/view?usp=sharing
122. Zurawski, J., "EPOC Support of the NSF CI Ecosystem", Invited Lightning Talk, NSF Cyberinfrastructure for Major Facilities Workshop, March 1, 2022. Slides available at: <https://drive.google.com/file/d/1aGbLmoQXcS2IPrRfg9aiSZzVOAPJbNWh/view?usp=sharing>
123. Zurawski, J., "EPOC Overview", Invited Lightning Talk, NOAA Stakeholders meeting, March 3, 2022.
124. Zurawski, J., Miller, K., "Engagement and Performance Operations Center (EPOC) Overview and Functions" Invited Talk, Pennsylvania State University Research State of the Practice Group, March 3, 2022. Slides available at: https://drive.google.com/file/d/1QL_iYI2twfx92-sbd117hJgEY6HqKo5K/view?usp=sharing
125. Southworth, D., "NetSage: Flexible Monitoring Framework for Multiple Applications", Invited Talk, GÉANT 2nd Performance Management Workshop, March 8, 2022. Slides available at: https://drive.google.com/file/d/1CihIkBNuZVd4W4B2v_EspouEqJZTZ-p9/view?usp=sharing
126. Miller, K., "Engagement and Performance Operations Center: Data Mobility Exhibition (DME)", Invited Talk, Quilt Spring meeting, March 10, 2022. Slides available at: <https://drive.google.com/file/d/1K0uX3Wl4gHBrVE8iQ4oP0ZMdvEStKzOL/view?usp=sharing>
127. Schopf, J., "Engagement and Performance Operations Center (EPOC) Now and Next", Invited Talk, Texas Advanced Computing Center (TACC), March 15, 2022. Slides

available at: <https://docs.google.com/presentation/d/1FWyQXyyHCtthC0fc9fqaF8-0toAkWN4k2AuTIX-nvE8/edit?usp=sharing>.

128. Meade, B., Addleman, H., and Mitchell, W., “GNA-G Routing Working Group Update”, Invited Talk, Routing Group, March 30, 2022. Slides available at: <https://docs.google.com/presentation/d/1qCQn6wF3UaGmU9rTgg8swOCFY4WpUTHaD2hLx3J6mUo/edit?usp=sharing>.

13. Regional Partner Activities

EPOC is partnered with the eleven Regional Partners:

- **Front Range GigaPop (FRGP)** is the regional collaboration of networks that cover the western states, including Colorado, Wyoming, Arizona, Idaho, Utah, and New Mexico.
 - *Roadside Assistance and Consultation Cases*: 63, 76, 166, 169, 198, 222
 - *NetSage*: <http://frgp.NetSage.global>, deployment February 2020, and discussions are ongoing for an SNMP deployment. Also, NetSage Tstat deployment at the NCAR Wyoming Data Center.
 - *Planning*: As part of our Partner Coordination Meeting, we agreed to give an overview talk at an upcoming technical meeting. However, scheduling this has been complicated and we hope to do this in the next quarter.

- **The Great Plains Network (GPN)** is the regional network that serves North Dakota, South Dakota, Nebraska, Iowa, Minnesota, Kansas, Missouri, and Arkansas.
 - *Meetings*:
 - Zurawski attended the 2021 KanREN Annual Meeting, April 21-22, 2021 [10], and Schopf, Addleman and Southworth attended the GPN 2021 Virtual Annual Meeting, June 2-4, 2021 [13], giving multiple presentations
 - *Roadside Assistance and Consultation Cases*: 153, 177, 178, 194, 202
 - *Deep Dives*:
 - Virtual Deep Dive with University of South Dakota was completed [108]
 - St Louis University (SLU) has participated in a Virtual Deep Dive, but the report is pending
 - *NetSage*: <http://gpn.NetSage.global>, SNMP and Flow data
 - *DME*: Members of GPN participated in the DME
 - *Training*: Meade, Addleman, and Zurawski attended the virtual Programmable Switches Workshop, virtual, February 16-17 and 23, 2022, http://ce.sc.edu/cyberinfra/p4_workshop_feb_2022.html, sponsored by GPN, University of South Carolina, NSF, and EPOC [54]
 - *Planning*: At our earlier Partner Coordination Meeting, we agreed to participate in the June 2022 GPN All Hands Meeting and are jointly making plans for this work.

- **iLight** is the regional network for Indiana.
 - *Meetings*:
 - Addleman attended and presented at the I-Light Town Hall Meeting on August 31, 2021 [28].
 - Moynihan and Meade attended the I-Light Members Meeting on September 28-29, 2021, <https://ilight.net/members-meeting/agenda/> [33].
 - *Roadside Assistance and Consultation Cases*: 172, 181, 188, 209, 213
 - *NetSage*: <http://ilight.NetSage.global>, deployment May 2019

- **The Keystone Initiative for Network Based Education and Research (KINBER)** is the regional network for Pennsylvania.

- *Meetings:* Zurawski and Miller attended a meeting with Pennsylvania State University, March 3, 2022. They gave an EPOC overview for their monthly Research State of the Practice Group. [58]
 - *Roadside Assistance and Consultation Cases:* 105, 182, 183, 184, 218, 225
 - *NetSage:* <http://kinber.NetSage.global> Deployment November 2020, but due to changing hardware and equipment ownership, this was discontinued in Year 5 Q4.
 - *DME:* Members of KINBER participated in the DME.
- **The Lonestar Education and Research Network (LEARN)** is the regional network for Texas.
 - *Meetings:*
 - Southworth attended the LEARN Technical Advisory Group (TAG) Meeting, June 8-10, 2021 [14]
 - Zurawski attended the LEARN Board meeting on December 7, 2021, and answered questions related to EPOC Deep Dives supporting the region [48].
 - *Roadside Assistance and Consultation Cases:* 113, 155, 162, 167, 170, 176, 180, 185, 193, 217, 227
 - *Deep Dives:*
 - A virtual Deep Dive took place with South Plains College, jointly with LEARN, November 19, 2021. A report is expected in mid 2022.
 - An in-person Deep Dive took place with McLennan Community College, jointly with LEARN, in March 2022. A report is expected in mid 2022.
 - An in-person meeting is being planned for April 2022 with LEARN and Midland, with a report is expected in mid 2022.
 - *NetSage:* Initial deployment of Flow data completed
 - *Training:* Southworth, D., “perfSONAR Training”, LEARN Technical Advisory Group (TAG) Meeting, June 8 & 10, 2021, Slides available online at: <https://drive.google.com/drive/folders/1ZSgwlrzwfZ7gS-NBtLrEc7ejRSI98M-y?usp=sharing> [81]
- **NOAA N-Wave**, the R&E network for NOAA
 - *Meetings:*
 - Zurawski attended and presented virtually at the NOAA N-Wave Joint Engineering and Technical Interchange (JETI) Conference, August 4-6, 2021, <https://www.noaa.gov/organization/information-technology/n-wave-joint-engineering-and-technical-interchange-jeti> [23]
 - Zurawski attended and presented at the NOAA Virtual Stakeholders Meeting, March 1-3, 2022 [55]
 - *Deep Dives:*
 - A virtual Deep Dive took place with the NOAA Marine Mammal Acoustics (NOAA NMFS) in summer 2021, and a final presentation occurred at NOAA N-Wave Joint Engineering and Technical Interchange (JETI)

conference. The effort was completed with a final report in August 2021 [91].

- **The Ohio Academic Resources Network (OARnet)** is the regional network for Ohio.
 - *Roadside Assistance and Consultation Cases*: 174, 179, 186
 - *Deep Dives*: A virtual meeting was held in December 2021 with OARnet and Sinclair staff, and a report is expected in mid 2022.
 - *NetSage*: A NetSage deployment will not be feasible until the R&E network is restructured.

- **Pacific Islands Research and Education Network (PIREN)** and the University of Hawai'i System Network, provide R&E network capacity to interconnect Pacific Islands with each other and the global R&E network fabric, including links to Australia and Guam, in addition to connectivity for the University of Hawai'i system and Mauna Kea and Haleakala astronomy observatories.
 - *New Member*: Joined Year 4 Quarter 3.
 - *Roadside Assistance*: 121, 208, 211
 - *NetSage*: Flow and SNMP now hosted by EPOC, started with previous IRNC support.

- **Pacific Northwest GigaPop (PNWGP)** provides access to next generation internet services and technologies throughout the Pacific Rim, but in the US primarily in California, Oregon, and Washington State.
 - *Roadside Assistance and Consultation Cases*: 117, 142, 189, 212
 - *Deep Dive*: EPOC is partnering with San Diego State University to perform a Deep Dive of campus drivers. The goal is to perform this as a hybrid event, with virtual questions that started in February of 2022, and an in-person event scheduled for April of 2022. A report is expected in Summer of 2022.
 - *NetSage*: Flow and SNMP now hosted by EPOC, started with previous IRNC support.
 - *Training*: CENIC training event, joint with U of SC, being discussed for Fall 2022.

- **Southern Crossroads (SoX)** is the regional network for much of the southeastern part of the US, including parts of Alabama, Georgia, South Carolina, and Tennessee
 - *Meetings*:
 - Miller attended and presented at the Fall SoX Participant Meeting, October 7, 2021, <https://www.sox.net/2021-fall-participants-meeting/> [34].
 - Addleman and Meade, with Cas D'Angelo from SoX, jointly facilitated a breakout session at the NSF CC* PI Workshop on September 14-15 and 21-22, 2021, <https://www.thequilt.net/public-event/2021-nsf-virtual-cc-pi-workshop/> [29].
 - Southworth participated in a virtual SoX Engineering meeting, February 24, 2022. He presented a tutorial on NetSage and EPOC roadside assistance. [53]

- *Roadside Assistance and Consultation Cases*: 168, 214, 233
- *NetSage*: Deployment of NetSage for Flow data went live in July 2020, in discussions for SNMP deployment.
- *DME*: Members of SoX participated in the DME.

- **Sun Corridor Network (SCN)** is the regional network for the state of Arizona
 - *Meetings*: Schopf presented at the Sun Corridor Network Engineering Meeting, June 14, 2021. She gave a presentation on NetSage and how it could be used by Sun Corridor and its members.
 - *Roadside Assistance and Consultation Cases*: 143, 160, 190, 191
 - *Deep Dive*: A Virtual Deep Dive took place January 10-12, 2022, for Arizona State University and Sun Corridor. A report is in progress which we hope to publish in early Year 5
 - *NetSage*: An initial NetSage deployment for Sun Corridor was made public in March 2021 at <https://suncorridor.netsage.global/>, but a network re-configuration started in early 2022 and all data collection efforts have been temporarily paused.

- **Texas Advanced Computing Center (TACC)** at the University of Texas at Austin, United States, is an advanced computing research center
 - *Meetings*: Schopf visited TACC in March 2022 and gave a talk discussing EPOC futures [127].
 - *Roadside Assistance and Consultation Cases*: 107, 113, 135, 227
 - *NetSage*: Deployment for TACC completed in July 2020. Archive deployment in production since 2019.