

Engagement and Performance Operations Center (EPOC)

(Formerly known as ReSEC)

NSF Award #1826994

Year 3 Quarter 4 and Annual

1 April 2020 through 31 March 2021

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/Portal, the provision of Managed Services, and Training. Year 3 highlights include adding SoX, Sun Corridor Network, and CENIC and Regional Networking Partners, work on 97 Roadside Assistance and Consultation Cases, three new NetSage deployments, and two new Deep Dive reports published.

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 to test transfer times against known “good” end points, with the goal of transferring a Terabyte in an Hour;
- Provision of Managed Services via support through the IU Global NOC and our Regional Network 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 3 as well as plans for Year 4. Note that at the time of this report, COVID-19 and its associated travel restrictions were in a state of high fluctuation.

2. Staffing and Internal Coordination

At the beginning of Year 3, 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
- Scott Chevalier, IU, Science Engagement and Training
- Dan Doyle, IU, system architect - Measurement and Monitoring co-Lead
- Heather Hubbard, IU, Project coordination
- Ken Miller, LBNL, Science Engagement and Managed Services support
- Ed Moynihan, IU, Science Engagement
- George Robb, LBNL, Managed Services support
- Doug Southworth, IU, Partner coordination and Deep Dive support

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.

Scott Chevalier was laid off in December 2020 due to a lack of project funding for perfSONAR focused work. IU has posted a position for another network engineer in February, and plans to interview candidates in April. The university is currently in a hiring freeze so the process is taking longer than one would hope.

The EPOC staff coordinate internally via four primary mechanisms:

- Synchronous and asynchronous communication via an email mailing list and Slack workspace;
- Project management via shared Trello (digital KanBan board) to track ongoing projects, requests, and record outcomes;
- Weekly project team meetings to update the Trello infrastructure and triage new requests; and
- Twice yearly All Hands Meetings for face-to-face discussion on important strategic topics.

Instead of a formal All Hands Meeting (AHM) this project year, we held a series of 2-3 hour Focus Times, each on an area of concentration that needed more time or discussion than the weekly meetings would allow. These have been relatively successful in moving forward longer-term issues, such as developing the virtual Deep Dive process, evaluating next steps for BGP, and identifying proactive opportunities for Consultations. These will likely continue over the next project year.

3. 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 first three quarters of Year 3, these included:

- Doyle and Schopf attended the GlobalNOC User Meeting, April 13-14, 2020, <https://usergroup.globalnoc.iu.edu/events/globalnoc-days-2020/>.
- Addleman presented at the Large Scale Networking (LSN) Workshop on Huge Data: A Computing, Networking and Distributed Systems Perspective, April 13-14, 2020, <https://protocols.netlab.uky.edu/~hugedata2020/>.
- Addleman attended the FABRIC Community Workshop, April 15-16, 2020, <https://whatisfabric.net/events/fabric-community-workshop-2020>.
- Moynihan attended the Science Gateways Community Institute (SGCI) partners meeting, April 21, 2020.
- Zurawski attended the GlobusWorld, April 29, 2020, <https://www.globusworld.org/conf/>.
- Addleman, Zurawski, Chevalier, Southworth, Miller, and Eli Dart attended and presented at the Training Workshop for Educators and Network Engineers on High Speed Network Protocols and Security, May 4, 2020, http://ce.sc.edu/cyberinfra/workshop_2020.html.
- Zurawski attended the LEARN Member Meeting, June 1-3, 2020.
- Southworth presented at the WestNet Meeting, June 24-26, 2020, <https://www.westnet.net/new-events>.
- Schopf attended the PEARC'20 meeting, July 27-31, <https://pearc.acm.org/pearc20/>.
- Moynihan attended the TICAL 2020 conference, August 31-September 4, 2020, <https://tical2020.redclara.net/>.
- Schopf and Addleman attended the GNA-G Meeting, September 14-15, 2020.
- Moynihan attended the LHCONe meeting, September 16-17, 2020 <https://indico.cern.ch/event/932306/>.
- Southworth and Chevalier attended and presented at the LEARN perfSONAR Virtual Training, September 21-22, 2020.
- Addleman attended the NSF Virtual CyberSecurity Summit, September 22-24, 2020, <https://www.trustedci.org/2020-nsf-summit>.
- Schopf, Addleman, Miller, and Zurawski attended the Quilt Fall Virtual Meeting, September 30 - October 1, 2020, <https://www.thequilt.net/public-event/2020-quilt-virtual-fall-member-meeting/>.

- Schopf and Addleman attended the Internet2 TechExtra meeting, October 6-7, 2020, <https://www.internet2.edu/news-events/events/techextra-2020/>.
- Schopf attended and presented at the Southern Crossroads (SoX) All Hands Meeting, October 8, 2020.
- Schopf attended the Fall CASC meeting, October 13-16, 2020, <https://casc.org/event/casc-fall-2020-membership-meeting/>.
- Addleman attended the Open Science Grid (OSG) Campus Workshop, October 22-23, 2020, <https://indico.fnal.gov/event/45998/overview>.
- Schopf attended the EduCause Annual Conference, October 27-29, 2020, <https://events.educause.edu/annual-conference>.
- Addleman attended the NSF Large Scale Facilities workshop, October 28, 2020, <https://largefacilitiesworkshop.com/>.
- Chevalier and Southworth attended the Internet2 perfSONAR Day, November 2, 2020, <https://internet2.edu/past-events/perfsonar-day-2020-events/#perfsonar>.
- Southworth attended the I-Light All Hands Meeting, November 4-5, 2020, <https://ilight.net/members-meeting/>.
- Southworth attended the GÉANT Telemetry meeting, November 10, 2020, <https://wiki.geant.org/display/PUB/Telemetry+and+Big+Data+Workshop>.
- Addleman, Schopf, and Zurawski attended the SC20 Conference, November 9 -19, 2020, <https://sc20.supercomputing.org>.
- Schopf attended the 7th International Workshop on Innovating the Network for Data Intensive Science (INDIS 2020) Workshop, part of SC'20, November 12-13, 2020, <https://scinet.supercomputing.org/community/indis/previous-editions/sc20-indis/>.
- Addleman attended the Workshop on Programmable Networking, November 16-18, 2020, <https://sites.google.com/view/us-japan-workshop/home>.
- Schopf and Zurawski attended the NOAA Joint Technical Interchange (JTI) meeting, December 1-3, 2020, <https://www.noaa.gov/organization/information-technology/n-wave-joint-engineering-and-technical-interchange-jeti>.
- Addleman attended the Fabric Partners Workshop, December 3-4, 2020, <https://fabric-testbed.net/events/facility-partner-workshop-2020>.
- Schopf and Addleman attended the GNA-G Meeting, December 7-8, 2020.
- Miller and Southworth attend the KINBER All Hands Meeting, December 8, 2020.
- Addleman, Miller, Robb, and Southworth attended the GÉANT DTN Infoshare, December 9, 2020, video link: https://youtu.be/-DSq_LCw36k.

Remote participation in Quarter 4 included:

- Schopf participated in NSF Workshop on Overcoming Measurement Barriers to Internet Research (WOMBIR 2021), January 11-12, 2021, <https://www.caida.org/workshops/wombir/2101/>. Her white paper on needing additional data for routing was accepted, and she participated in discussions around internet measurements and use cases.
- Schopf, Addleman and Moynihan attended the Annual TransPacific Research and Education (TPRE) Networking Meeting, January 16, 2021. They discussed networking updates and engineering issues.

- Schopf attended the ESIP Winter meeting, January 26-29, 2021, <https://2021esipwintermeeting.sched.com/>. She went to several sessions to discuss data transfer use cases.
- Schopf attended and presented at APAN 51, February 1-5, 2021, <https://apan51.apan.net/>. She participated in several sessions and presented on the topic of network routing issues.
- Zurawski and Schopf attended the Quilt Virtual Winter Meeting, February 9 - 11, 2021, <https://www.thequilt.net/public-event/2021-quilt-virtual-winter-member-meeting/>.
- Zurawski attended the ERN Structural Biology Voice of the Customer Virtual Workshop February 11, 2021, <https://www.ernrp.org/event/ern-structural-biology-voice-of-the-customer-virtual-workshop/>.
- Schopf and Addleman FAB Kickoff Meeting, February 18, 2021. They participated in discussions on how to support FABRIC experimentation.
- Zurawski and Schopf attended the NOAA 2021 N-Wave Stakeholders Summit, February 23-25, 2021, <https://noc.nwave.noaa.gov/nwave/public/events/2019-n-wave-stakeholders-and-science-engagement-summit>. They presented concerning data needs for EPOC Deep Dives, and how these may be integrated into NOAA programs.
- Schopf participated in the Science Gateways sponsored "Jumpstart Your Sustainability Plan" mini-course, March 1-3, 2021, <https://sciencegateways.org/engage/focus-week/jumpstart>. The training session included tips and approaches to understanding and evaluating project sustainability.
- Schopf participated in both sessions of the GNA-G Spring Update meeting, March 15-16, 2021, <https://www.gna-g.net/attend-a-meeting/gna-g-community-vcs-2021q1/>. She and Addleman gave a presentation on the new Routing Working Group.
- Southworth and Addleman attended the The African Peering & Interconnection Forum (AfPIF) meeting "The Death of Transit and the Evolving Role of the IXP", March 23, 2021, <https://www.afpif.org/virtual-peering-series-africa/death-of-transit-the-evolving-role-of-ixps/>. They participated in discussions on content providers buying their own dark fiber and laying new subsea cables to bypass the traditional IP transit providers.
- Zurawski and Moynihan attended the LHCONE meeting, March 23-24, 2021, <https://indico.cern.ch/event/983436/>.

4. Presentations and Publications

For Year 3, the EPOC presentations, invited posters, and publications are listed here and referred to throughout the report with the reference number listed:

1. Addleman, H., "EPOC: Roadside Assistance to the Rescue!", Invited Talk, Large Scale Networking (LSN) Workshop on Huge Data: A Computing, Networking and Distributed Systems Perspective, April 13, 2020. Slides available online at: <https://epoc.global/wp-content/uploads/2020/10/20200403-EPOC-Roadside-Assistance-to-the-Rescue-BigData-Workshop.pdf>
2. Addleman, H., Dart, E., "Network Strategy to Enable Data Intensive Science", Invited Talk, CI Engineering Lunch and Learn, May 1, 2020. Slides available online at:

- https://drive.google.com/file/d/16vE-kmktR3YGFJYoYfZSTqU_rxv04Did/view?usp=sharing.
3. Zurawski, J., "Science DMZ and TCP", Invited Talk, Training Workshop for Educators and Network Engineers on High Speed Network Protocols and Security, May 4, 2020, Slides available online at:
http://ce.sc.edu/cyberinfra/docs/workshop/workshop_2020_online/TCP_DMZ-Zurawski.pdf.
 4. Chevalier, S., Southworth, D., "perfSONAR Overview", Invited Talk, Training Workshop for Educators and Network Engineers on High Speed Network Protocols and Security, May 4, 2020. Slides available online at: <https://drive.google.com/file/d/1AOxRnIVWi-4klqUFf14Mo42sZS9ep5SV/view?usp=sharing>.
 5. Addleman, H., Dart, E., "BGP Architectures and Best Practices", Invited Talk, Training Workshop for Educators and Network Engineers on High Speed Network Protocols and Security, May 5, 2020. Slides available online at:
<https://drive.google.com/file/d/1SomS5T6xB4Smko7Y-IB9BtnbtYeNaCNN/view?usp=sharing>.
 6. Miller, K., "Security best practices in high-speed networks", Invited Panel, Training Workshop for Educators and Network Engineers on High Speed Network Protocols and Security, May 6, 2020. Slides available online at: <https://drive.google.com/file/d/1VZ-RUOAtKik4GwM2zu7cw3ZoYKboo9Pq/view?usp=sharing>.
 7. Zurawski, J., Schopf, J., Addleman, H., "University of Wisconsin-Madison Campus-Wide Deep Dive", *Lawrence Berkeley National Laboratory Technical Report*, #2001325, May 26, 2020. Report available online at: <https://escholarship.org/uc/item/99c4v5xh>
 8. Zurawski, J., Schopf, J., "Deep Science Understanding Leads to Enhanced CI Capabilities", Invited Article, The Quilt Circle, June 5, 2020. Article available online at: https://issuu.com/noahredman1/docs/thequiltcircle_2020_5-29-20_withcover
 9. Southworth, D., "EPOC Overview, Cases, and Tools", Invited Talk, WestNet User Meeting, June 26, 2020. Slides available online at:
https://drive.google.com/file/d/1eDdHg6jwXfa-Elk4EEW5Lsze1WoDVQx_/view?usp=sharing.
 10. EPOC Newsletter #1, July 10, 2020. Newsletter available online at:
<https://epoc.global/wp-content/uploads/2020/08/EPOC-NL-7-10-2020.pdf>
 11. Southworth, D., "Whose Line is it Anyway?: Problem Solving in Complex Networks", Trusted CI Monthly Webinar Series, July 20, 2020. Video available online at:
<https://www.youtube.com/watch?v=TKH-6CNO2mc>
 12. EPOC Newsletter #2, August 7, 2020. Newsletter available online at:
<https://epoc.global/wp-content/uploads/2020/08/EPOC-NL-8-7-2020.pdf>
 13. Southworth, Douglas, "Introduction to perfSONAR", LEARN Technical Advisory Group webinar, August 10, 2020. Slides and video available at:
https://drive.google.com/drive/folders/1jMAG6uw_-ESvATcb895xg-JWc6b7oxjk?usp=sharing
 14. Zurawski, J., "Identifying Campus Infrastructure Needs", Quilt CC* Seminar, August 13, 2020. Video online at <https://youtu.be/EqNvKQWwYuQ>

15. Miller, K., "Data Movement Exhibition" Invited Talk, CI Engineering Lunch and Learn, August 14, 2020. Video available online at: <https://youtu.be/CmHGu9cG0ww>
16. Robb III, G., Pejovic, G., Shiflett, C., "Data Transfer Node Design Patterns" Invited Talk, CI Engineering Lunch and Learn, August 28, 2020, Slides available online at: https://drive.google.com/file/d/1ABSe4OalAROTBi0ItM5VXcLORQ2yaNa_/view?usp=sharing
17. Schopf, J., "NSF EPOC and University of Cincinnati Campus-Wide Deep Dive", Invited Talk joint with J. Combs, University Cincinnati, CaRCC Data Facing Track Call, September 1, 2020. Video available online at https://youtu.be/jhg5R9_VdQY
18. Miller, K., "Data Movement Exhibition", Invited Talk, Quilt Managed Services Working Group, September 8, 2020. Video online at <https://youtu.be/wuZrP2Rg1IA>
19. Southworth, D., and Chevalier, S., "perfSONAR Workshop", LEARN Virtual Training, September 21-22, 2020. Slides available online at: https://drive.google.com/drive/folders/1jMAG6uw_-ESvATcb895xg-JWc6b7oxjk
20. EPOC Newsletter #3, September 24, 2020. Newsletter available online at: <https://epoc.global/wp-content/uploads/2020/12/EPOCNLsept.pdf>
21. Schopf, J. "NetSage Now", invited presentation, CI Engineering Lunch and Learn, September 25, 2021. Video available online at: https://www.youtube.com/watch?v=9e4vK_reT_Y
22. Schopf, J., "Moving Data Faster with the Engagement and Performance Operations Center (EPOC)", Invited Talk, Quilt Fall Virtual Meeting, September 30, 2020. Slides available online at: <https://drive.google.com/file/d/1OshtyJniXVgRkPJP8apkP3jXFKMpl0M3/view?usp=sharing>
23. Miller, K., "Data Movement Exhibition", Invited Talk, Quilt Fall Virtual Meeting, September 30, 2020. Slides available online at: <https://drive.google.com/file/d/1NywUeYH4xvImiwb4pjqBvpxl9sPafTkl/view?usp=sharing>
24. Schopf, J., "NetSage, EPOC, and SoX", Invited Talk, SoX All Hands Meeting, October 8, 2020.
25. Schopf, J., Zurawski, J., "Learning about Researcher Cyberinfrastructure Needs with EPOC", Invited Talk, EduCause Annual Conference, October 27-29, 2020. Video available online at: <https://drive.google.com/file/d/1s2po49sNHf2wzjahkURQggShMc-tchm6/view?usp=sharing>
26. Southworth, D., "IRNC NetSage - Use Cases and Scalability", Invited Talk, GÉANT Telemetry and Big Data Workshop, November 10, 2020. Slides available online at: <https://drive.google.com/file/d/1DYaBxuC5TnzWynvbtn2a057QX6Vg0ytA/view>
27. K. Turner, E. Balas, S. Baveja, D. Doyle, L. Ensman, S. Faci, A. Gonzalez, M. Khanal, A. Lake, J. Leigh, T. Seto-Mook, D. Southworth, B. Tierney, and J. M. Schopf, "The NetSage Measurement Framework: Design, Development, and Discoveries", *Proceedings of INDIS'20*, Atlanta, Georgia, USA, November 16, 2020.
28. Zurawski, J., "Institutional of Data Needs: EPOC Deep Dives", Invited Talk, NOAA Joint Technical Interchange (JTI) meeting, December 1-3, 2020. Slides available online at:

- <https://drive.google.com/file/d/1v-2YBAR7IDM2t1H-INSrvaCuQeetTb20/view?usp=sharing>
29. Zurawski, J., "I think I need a 100G DTN! EPOC Evaluation", EPOC Contact Point: Jason Zurawski, December 4, 2020. Available online at: <https://epoc.global/wp-content/uploads/2020/12/I-think-I-need-a-100G-DTN.2020-12-4.pdf>
 30. Zurawski, J., "I think I need a 100G perfSONAR Node! EPOC Evaluation", EPOC Contact Point: Jason Zurawski, December 7, 2020. Available online at: <https://epoc.global/wp-content/uploads/2020/12/2020-12-07-I-think-i-need-a-100G-perfSONAR-node-2.pdf>
 31. Miller, K., Southworth, D., "EPOC Overview for KINBER", Invited Talk, KINBER All Hands Meeting, December 8, 2020. Slides available at: <https://drive.google.com/file/d/1v0X1grAMynfXvvUqX0h46VhTKBT56bHS/view?usp=sharing>
 32. EPOC Newsletter #4, December 9, 2020. Newsletter available online at: <https://epoc.global/wp-content/uploads/2020/12/EPNL-Dec-9.pdf>
 33. Schopf, J., "EPOC support for the NSF CC* Program", invited webinar, The Quilt Network Webinar, December 15, 2020. Slides available online at: <https://drive.google.com/file/d/1lfKplQfD1RFJr5i-OULnHm9cCoeo1mcz/view?usp=sharing>
 34. Schopf, J., "EPOC support for the NSF CC* Program", Invited Talk, CaRCC Emerging Centers Track, December 16, 2020. Slides available online at: <https://drive.google.com/file/d/1y7YIkG8UN9DwLg3OzaE2thmO8nZbyad4/view?usp=sharing>
 35. Schopf, J., "NetSage, EPOC and OARnet", OARnet Engineering Meeting, January 6, 2021.
 36. Schopf, J., Tierney, B., Addleman, H., Southworth, D., "Understanding Routing Behaviors on R&E Networks", White Paper, *Proceedings of the NSF Workshop on Overcoming Measurement Barriers to Internet Research (WOMBIR 2021)*, January 11, 2021. Available online at: <https://drive.google.com/file/d/1WUOpy34ypbMB9Ro11zgYyceYdoZQwv5o/view?usp=sharing>
 37. Schopf, J., "EPOC support for the NSF CC* Program", Invited talk, CI Engineers Lunch and Learn, January 15, 2021. Video available online at: <https://www.youtube.com/watch?v=IRmMgyWUgcY>
 38. Schopf, J., "TransPAC5 and NEA³R", Invited Talk, Annual TransPacific Research and Education (TPRE) Networking meeting, January 16, 2021. Slides available online at: <https://docs.google.com/presentation/d/1kFSmx3opVpzfWa8Fm1JPY42wikLjVVnVcFaZzXAk8tE/edit?usp=sharing>
 39. Zurawski, J., Schopf, J. "Baylor University Campus-Wide Deep Dive". *Lawrence Berkeley National Laboratory Technical Report, #2001384*, January 21, 2021. Report available online at: <https://escholarship.org/uc/item/5x9234mv>
 40. Schopf, J., "Routing Problems and How to Address Them", BOF Talk, APAN 51, February 2, 2021. Slides available online at:

https://docs.google.com/presentation/d/1_I7ytZR_KmO_dpwFg1Bq0SO44XySoTlr8BEKc6zQnTY/edit?usp=sharing

41. Schopf, J., "TransPAC4 and TransPAC5 Update", Invited Talk, APAN 51, February 4, 2021. Slides available online at:
https://docs.google.com/presentation/d/1EUrPQ_7j40okohptev_ark_AW_LJcN2PIkRUK-joTak/edit?usp=sharing
42. Schopf, J., Lee, F., "Routing Problems and How to Address Them (Summary)", Invited Talk, APAN 51, February 5, 2021. Slides available online at:
https://docs.google.com/presentation/d/1GV7mLPTbJtPd0Vr2o3GXMw3_bEd48k3aOVI2N1xBUU/edit?usp=sharing
43. Schopf, J., "LEARN and NetSage", Invited Talk, LEARN Engineers Meeting, February 8, 2021. Slides available online at: https://drive.google.com/file/d/1uBzgKuPqRmlYZ-Svb9zCJ966Hr6_dNFP/view?usp=sharing
44. Schopf, Jennifer, "International Networks at Indiana University Support for FAB", Invited presentation for FAB kickoff meeting, February 18, 2021, slides available at:
https://docs.google.com/presentation/d/1ixj5CRGsQ-ngmKsuzaQYJS7beSQV_gho/edit#slide=id.p1
45. Zurawski, J., "Institutional Examinations of Data Needs: EPOC Deep Dives", Invited Talk, NOAA 2021 N-Wave Stakeholders Summit, February 23, 2021. Slides available online at: https://drive.google.com/drive/folders/12Y7s41194axDsF1VJdulkGBrbc3p_jXF
46. Schopf, J., "NetSage, TransPAC, and SingaREN", Invited talk, SingaREN Monthly Meeting, March 9, 2021. Slides available online at:
<https://drive.google.com/file/d/1AkmxL5DyQ2ZgVTf6Jcy4G2vdcCqie0Qb/view?usp=sharing>
47. Schopf, J., "GNA-G Routing WG", Invited Talk, GNA-G meeting, March 15, 2021. Slides available online at:
<https://drive.google.com/file/d/1JeWqX5oheTsk6u6RWPG77mIYJbmJJIEy/view?usp=sharing>
48. Schopf, J., "NetSage, EPOC, and the Sun Corridor Network", Invited talk, Sun Corridor Network Engineers meeting, March 16, 2021. Slides available online at
<https://drive.google.com/file/d/1tJ0cFouSP9SMnBXXG085zfA9unnT7Pzu/view?usp=sharing>
49. Schopf, J., "NetSage, EPOC, and NYSERNet", Invited talk NYSERNet Engineering Meeting, March 22, 2021. Slides available online at
<https://drive.google.com/file/d/1FeT4Yml-Un2sXPNAmb9lcQUXNlo8vR8D/view?usp=sharing>

5. Project Coordination

EPOC has three types of partners: *Regional Networking 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*.

5.A Additional Regional Networking Partners

Over the last eighteen months, several institutions have approached EPOC about the possibility of becoming formal partners. While some EPOC activities are available to the broader community, such as Roadside Assistance and Consulting and Training, others are limited to partners or partners receive scheduling preference, including NetSage deployments and work with Managed Services. Complementary to this, the EPOC project has experienced cost savings due to the reduction in staff travel, so it was determined that we could consider expanding our formal partner set at this time.

New Regional Networking Partners in Year 3 included:

- Sun Corridor Network (SCN), the regional network for the state of Arizona,
- Southern Crossroads (SoX), the regional network for much of the southeastern part of the US including parts of Alabama, Georgia, South Carolina, and Tennessee,
- The Texas Advanced Computing Center (TACC) at the University of Texas at Austin, United States, an advanced computing research center that supports Texas institutions as well as having a national mission,
- CENIC/Pacific Wave, the not-for-profits that support R&E networking for the State of California, as well as additional connections for the west coast of the US and internationally.

We are currently in conversations with NYSERNet, the R&E network for the state of New York, to see if they would also like to join EPOC. The University of Hawaii's Pacific Islands Research and Education Network also began discussions in Quarter 4 with EPOC staff to understand if they were eligible to become an EPOC partner.

We are also in the process of discussions with N-Wave, the R&E network for NOAA, to see if they would like to be a partner. Schopf and Zurawski attended the NOAA Joint Technical Interchange (JTI) meeting, which was held virtually on December 1-3, 2020, and participated in sessions as well as presented concerning data needs from EPOC Deep Dives [28]. Another presentation was given at the NWave stakeholder meeting in February 2021 [45]. After the meeting, discussions regarding starting a NOAA Deep Dive project continued, with a target of starting to profile a scientific use case in early 2021. Newsletter 4 [20] also highlighted joint work with NWave.

5.B Regional Networking Partners

EPOC is partnered with the ten Regional Networking Partners:

- **CENIC** operates the California Research and Education Network (CalREN), the state R&E network for California.
 - Meetings: Several email conversations took place about CENIC/PacWave joining as an EPOC partner, which was finalized in January.
 - Roadside Assistance and Consultation Cases: 74, 90, 101, 104, 117, 142, 146
 - NetSage: Flow and SNMP from IRNC support
 - DME: CENIC member institutions participated in the DME

- We provided a Letter of Collaboration to Los Nettos, a CENIC member, for the CC* program
- We provided a Letter of Collaboration to Deelman at University of Southern California, a CENIC member, for the NSF CICOE program
- Other: Mentioned joint work in Newsletter 3 [20]
- **Front Range GigaPop (FRGP)** is the regional collaboration of networks that cover the western states, including Colorado, Wyoming, Arizona, Idaho, Utah, and New Mexico.
 - Meetings: We attended and presented at the June Member Meeting [9]
 - Roadside Assistance and Consultation Cases: 63, 76, 106, 110, 116
 - Deep Dive: Discussions are taking place with NOAA for a virtual Deep Dive
 - 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.
 - DME: FRGP member institutions participated in the DME
 - Other: Mentioned joint work in Newsletter 2 [12]
- **iLight** is the regional network for Indiana.
 - Meetings: Southworth attended the iLight All Hands Meeting in November.
 - NetSage: <http://ilight.netsage.global>, deployment May 2019, working on ASN splitting with SWIP.
 - We provided a Letter of Collaboration to Marru at IU, an iLight member, for the NSF CSSI program
- **The Great Plains Network (GPN)** is the regional network that serves North Dakota, South Dakota, Nebraska, Iowa, Minnesota, Kansas, Missouri, and Arkansas.
 - Meetings: We are planning to be part of the summer GPN All Hands Meeting
 - Roadside Assistance and Consultation Cases: 50, 59, 72, 86, 133, 148
 - Deep Dives: Virtual Deep Dive with University of South Dakota ongoing.
 - NetSage: <http://gpn.netsage.global>, SNMP and Flow data
 - DME: GPN member institutions participated in the DME
 - Managed Service: GPN members investigating use of Portal service, currently on hold.
 - We provided Letters of Collaboration to the University of Arkansas and OneNet, both for the CC* program
 - Other: Mentioned joint work in Newsletter 2 [12]
- **The Keystone Initiative for Network Based Education and Research (KINBER)** is the regional network for Pennsylvania.
 - Meetings: Miller and Southworth gave a presentation at the KINBER All Hands Meeting in December [31]
 - Roadside Assistance and Consultation Cases: 73, 87, 105, 130, 132
 - NetSage: <http://kinber.netsage.global> Deployment November 2020, working on ASN extensions with SWIP.
 - DME: KINBER member institutions participated in the DME
 - Managed Services: We finished work with Arcadia University and KINBER for a perfSONAR managed service at Arcadia.

- **The Lonestar Education and Research Network (LEARN)** is the regional network for Texas.
 - Meetings: Attended the June 2020 Member Meeting.
 - Roadside Assistance and Consultation Cases: 62, 69, 71, 113, 115, 123, 145, 155 (Note: we include these for both LEARN and TACC)
 - Deep Dives: Published report for Deep Dive with Baylor and LEARN that took place in January 2020 [39]; Discussions are taking place with TAMUSA for a virtual Deep Dive
 - NetSage: Continued discussion for deployment, including a talk at the LEARN Engineering Meeting [43]
 - DME: LEARN member institutions participated in the DME
 - Managed Service: LEARN-member Baylor University investigating use of Portal service, currently on hold
 - Training: We gave a talk on perfSONAR basics [13] and then led a hands-on training session in September 2020 [19]
 - We provided a Letter of Collaboration to LEARN for the CC* program
 - Other: Mentioned in Newsletter #4
- **The Ohio Academic Resources Network (OARnet)** is the regional network for Ohio.
 - Meetings: We met with OARNet leadership in lieu of their canceled AHM
 - Roadside Assistance and Consultation Cases: 35, 51, 57, 93, 111, 131, 144, 151
 - NetSage: A meeting and talk were given in January 2021 to discuss deployment options [15]. A deployment will not be feasible until the R&E network is restructured.
 - We provided a Letter of Collaboration to OARnet for the CC* program
 - Other: Joint work was mentioned in Newsletter 3 [20]
- **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: Schopf attended the SoX All Hands Meeting in October and gave a presentation on the NetSage installation [3]
 - Roadside Assistance and Consultation Cases: 88, 129
 - NetSage: Deployment of NetSage for Flow data went live in July 2020, in discussions for SNMP deployment.
 - DME: SoX member institutions participated in the DME
- **Sun Corridor Network (SCN)** is the regional network for the state of Arizona
 - Roadside Assistance and Consultation Cases: 81, 83, 99, 118, 141, 143, 160
 - NetSage: A presentation was given to the SCN Engineers meeting [48], leading up to a deployment that went live in March 2021.
 - Deep Dive: Discussions for a Deep Dive to take place in 2021 are ongoing.
 - DME: SCN member institutions participated in the DME
 - We provided Letters of Collaboration to Arizona State University for the CC* program

- **Texas Advanced Computing Center (TACC)** at the University of Texas at Austin, United States, is an advanced computing research center
 - Meetings: Discussions for membership took place this quarter, and they officially joined in June 2020.
 - Roadside Assistance and Consultation Cases: 62, 69, 71, 107, 113, 115, 123, 135, 145, 155 (Note: we include some of these for both LEARN and TACC)
 - NetSage: Deployment for TACC completed in July 2020. Archive deployment in production since 2019.
 - Deep Dive:
 - Managed Services:
 - Other: Joint work mentioned in Newsletter 4 [32]

Because many of the partners have had their yearly All Hands Meetings delayed or canceled, we plan additional check-ins with each partner to ensure progress on various deliverables.

5.C 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 gave a talk as part of the Data Facing Track on Deep Dives [17] and another for the Emerging Centers Track on EPOC support for the NSF CC* solicitation [34].
- **Trusted CI: The NSF Cybersecurity Center of Excellence** supports cybersecurity for NSF funded projects. Addleman is the main contact for Trusted CI. Southworth gave a talk for the Trusted CI Monthly Webinar series [11]. Addleman was interviewed by Trusted CI for part of their Success Stories blog series. This interview focused on the consultation between Trusted CI and International Networks at Indiana University. This consultation resulted in a security plan that is currently in use by both the NEA3R and TransPAC projects. Work with Trusted CI was also mentioned in Newsletter 1[10].
- **Internet2** supports solving common technology challenges for their over 200 educational, research, and community members. Schopf is the main contact for this organization. Several events typically held in person (The Internet2 Global Summit, Internet2 TechEx) were not held due to the pandemic, but a series of virtual replacements (Tech Expo, perfSONAR Training) were attended and led by EPOC team members. Joint work with Internet2 was also mentioned in Newsletter 4 [20].
- **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 in a number of ways, and has continued to do so virtually during the pandemic:

- Publication of EPOC materials for the Quilt Circle Magazine in 2020 [8], and a pending article for 2021
- EPOC staff are involved in the ongoing Quilt-run seminars related to the CC* program [14,18]
- EPOC staff participated in the Virtual Fall Quilt Member Meeting [22,23]
- Schopf also gave a webinar entitled “EPOC Support for NSF CC*” to the Quilt Community on December 15, 2020 [33]. This content focused on EPOC’s support of changes within the NSF CC* program, and ways EPOC could assist the regional community, and their members
- EPOC staff participated in the Virtual Winter Quilt Member Meeting
- **The Science Gateway Community Institute (SGCI)** provides best practice recommendations and support for scientists building and using data portals. In Quarter 1, Moynihan attended an internal SGCI partners video conference to learn more about their plans for 2020 and how they are adjusting to COVID-19 travel restrictions. In Quarter 2, SGCI updated their partnership program based on community input, and we affirmed our status as project partners for 2021. Schopf also participated in the Gateways Sustainability Workshop in March 2021.
- **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 updating XSEDE liaisons during the quarter. She was elected as a member of the XSEDE Advisory board and participated in PEARC, formally the XSEDE meeting.

Although not a formal Infrastructure Partner, we also worked closely with the Globus Team for several consultations this year, including #75, 109, 114, 120, 122, 126, 127, and 159. In addition, we are working closely with the Globus team on the test sites for the DME (Section 9.A) and participated in GlobusWorld 2020.

Another project we are in discussions with to see if they would like to formally become an Infrastructure Partner is with the FABRIC testbed, and its international component, FAB. Schopf and Addleman attended and presented at a number of FABRIC and FAB workshops and meetings, including the April 2020 Community Workshop, the December 2020 Partners workshop, and the FAB Kickoff meeting, where we gave a presentation as well [44]. We will be supporting several of the experiments and making sure the end-to-end paths for the testbed are running smoothly. FABRIC plans to start bringing experiments between institutions online later this calendar year, and EPOC engineers will work to make sure performance is as expected.

5.D 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. However, we have found out that most of the small to medium sized teams on the campuses we have worked with to date are not heavily involved in these efforts. Most teams are working independently. Because of this, we are decreasing our coordination with several of the less responsive Science

Community Partners, and instead will focus more strongly on the science collaborations identified through Consultations and Deep Dives.

The Year 3 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 attended the virtual Summer meeting, originally to be held in Burlington, VT, and the virtual annual Winter meeting in January 2021. Discussions continue with ESIP leadership on the possibility of doing an EPOC Deep Dive at an upcoming meeting and what that might look like given the move to all virtual meetings.
- **The University of Hawai'i System Astronomy Community** supports 15 facilities with hundreds of researchers and experiments every year. Data was added to the NetSage science registry with this project, and staff members are trailing the Data Portal. We are also coordinating on Consultations #57 and #119.
- **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, a conference which includes all of the regional hubs, in June 2020, but this meeting was cancelled due to COVID related travel restrictions and no updated meeting date has been set at this time.

5.E Advisory Committee

The EPOC External Advisory Board (EAB) members met virtually on April 17 2020 and were asked to offer feedback on key project goals, including:

- Virtual meetings (training, lessons learned, etc.)
- Developing services that are 'low touch', but still serve a scientific/research community
- Challenges they are facing, and ways EPOC services could be adapted to address

The EAB provided valuable feedback that indicated that EPOC should not attempt to dilute core approaches, such as Deep Dives, just to adapt to an opportunistic virtual environment. Other items, such as targeted training, were worth pursuing if the outcome would lead to advancement of EPOC mission space. Most, if not all, members agreed that there will be a fundamental shift in a post-COVID-19 environment where R&E sites may have to make hard choices regarding supporting research activities.

5.F External Partners

In addition to the previously mentioned partners, the EPOC team is coordinating with a number of 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. Current coordination is primarily taking place during the bi-weekly PRP/I2 Engagement calls, or when we jointly attend meetings.

We are working closely with members of the University of South Carolina Cyber Training team (MNSF#1829698), as detailed in Section 11.A.

EPOC is working with the CI Engineering Community (<https://www.es.net/science-engagement/ci-engineering-brownbag-series/>) to catalog the presented webinars from the Lunch and Learn series held from 2017 to present. To date, 129 webinars have been uploaded to the EPOC YouTube channel as of March 31, 2021, and available online at: <https://www.youtube.com/channel/UChlraulc1bccif1Dz4cfZI0w>.

During the project year, we expanded several informal partnerships in significant ways. We had previously been working with researchers at the Arecibo Observatory to transfer data from Puerto Rico to the continental US and advising on possible approaches. In December, the observatory suffered a catastrophic failure, and this work was raised in priority significantly. Working jointly with a team from University of Central Florida (UCF), the Texas Advanced Computing Center (TACC), the University of Puerto Rico, Globus, and the CCoE Pilot, we have begun discussion to shift the data to TACC. This is being tracked as Roadside Assistance Case #107.

In addition, in October a conversation started between various NSF IRNC network operators and the researchers who are part of the international Event Horizon Telescope (EHT) project. After initial discussions where the data transfer needs across the EHT partners were detailed, a smaller set broke off to try to make progress looking specifically at transfers between the France and Spain telescopes and the collector site at MIT. The initial steps have been to detail the end-to-end networking diagrams and to run basic tests for baseline numbers. This is being tracked as Roadside Assistance Case #121.

5.G Support for Grant Submissions

The NSF CC* Solicitation was released in December 2020. We gave three presentations [33, 34, 37] on how EPOC can support teams in both the proposal and award stage of a CC* project. We had fourteen consultations related to CC* proposals. We also reviewed and provided letters of collaboration to eight organizations:

- Arizona State University (ASU)
- Lonestar Education and Research Network (LEARN)
- Los Nettos
- Ohio Academic Resources Network (OARnet)
- OneNet
- University of Arkansas (UA)
- University of Mississippi (UMiss)
- University of North Carolina at Charlotte (UNC-C)

Separately from the CC* program, we provided a Letter of Collaboration for Suresh Marru, IU, for an NSF CSSI program, for a proposal entitled "FRAMEWORKS: Cybershuttle:

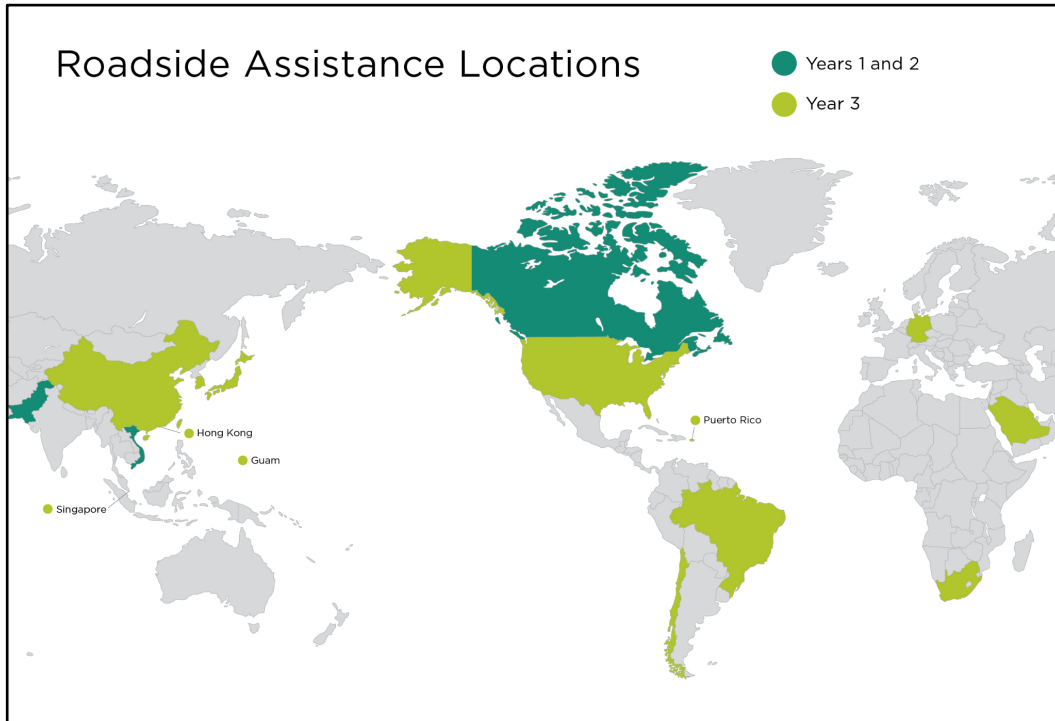


Figure 2: A map showing our Roadside Assistance and Consultations Cases internationally. Light green shows states where we had interactions with an institution in Year 3, and dark green shows a prior interaction only in Year 1 or 2.

In Year 3, we participated in 97 total Cases. Some additional statistics about the Cases include:

- We worked with institutions in 28 states and regions, including 21 EPSCoR jurisdictions, as shown in Figure 1.
- Cases with an international component included 12 different countries, as shown in Figure 2.
- 47 of 97 Cases, or 48%, were associated with a specific EPOC Regional Networking Partner.

6.A Roadside Assistance Cases

In Year 3, we had four completed Roadside Assistance Case:

- **59 - Saint Louis University (SLU):** A biologist at Saint Louis University (SLU) experienced very poor file transfer performance and reached out to EPOC. Transfers from his office to Amazon’s S3 cloud storage servers were so slow he was bringing disks home to upload on his home network to save time. Initial triage discovered that there were multiple firewalls in the transfer path for the biology data. In addition, some of the switches in the path had not been designed to support high-speed long-distance transfers and were also old enough to no longer be supported by the vendor. The biologist who initially reported the problems set up several perfSONAR nodes for continued testing, but then left SLU. EPOC staff discussed the issue with the SLU Director of IT services and concluded that it would be best to close this case since there was a full campus network refresh underway and the original reporter had left SLU.

- **71 - Texas Tech University (TTU):** TTU has 130 weather stations spread across the western US that are connected online in a variety of ways, including cellular modem, cable modem, and standard campus networks. They are each polled every 2.5 minutes by two separate servers on the TTU campus and each station sends back a few kilobyte data payload to both servers. On Thursday December 12, 2019, between twenty and fifty percent of the stations didn't respond between 11:50am-12:35pm and 2:50pm-3:35pm central time. This problem occurred every Thursday until February 13, 2020. The intermittent and complex nature of the problem made it very difficult to troubleshoot. The only change made shortly before February 13 was that the original collection server was retired and replaced by new hardware running the same collection program, which should not have affected this issue. Due to the intermittent nature of this problem, we held this engagement open for 10 weeks past the February problem resolution. TTU engineers reported that the problem had not re-occurred at that time.
- **95 - City College of New York (CCNY):** Engineers at the JGN reported packet loss across a tunnel between a JGN device at Kyutech Institute in Japan and a workstation endpoint at CCNY. EPOC staff engaged engineers at CCNY to set up perfSONAR tests between the sites. At that time, testing outside of the campus network had not replicated the packet loss. Several configuration changes took place during this case at CCNY and they engaged their campus engineers to look at possible firewall issues. EPOC staff spent time in Quarter 3 troubleshooting the on campus perfSONAR nodes and working with CCNY staff to run perfSONAR tests from various perfSONAR nodes located on R&E national networks in an effort to isolate the faulty equipment. COVID protocols on the CCNY campus delayed troubleshooting as well. A faulty device in the CCNY network was identified and replaced by campus engineers. JGN staff set up the GRE tunnel again and tested through the tunnel between the end hosts at CCNY and Kyutech. Testing results showed no packet loss between the hosts and greatly improved transfer rates. Results showed a 10x increase in speed from Kyutech to CCNY. With these results the tunnel was put back into production and researchers started data again. CCNY staff declined to share the make and model of the device being replaced or its replacement model.
- **128 - European Bioinformatics Institute (EBI):** As part of our efforts to actively investigate performance issues for data transfers related to COVID-19 research, we began looking at performance on increases in traffic between EBI and the NYGC. Initial investigation showed that the average rate between the two sites averaged ~20Mbps and that there was significant variance between hosts at NYGC with some hosts getting ~400Mbps and others only ~5Mbps. In November, we reached out to both EBI and NYGC in an attempt to gather more information about the end-to-end paths. After running perfSONAR tests and examining trace route information, we discussed the results with engineers from both EBI and NYGC and concluded that there were likely two possible causes for the performance variance, either the NYGC's use of a VPN for flows from EBI or the firewall and load balancer settings at EBI. At the end of November, EBI made internal adjustments to their security settings that improved average performance for these flows from ~20Mbps to ~100Mbps. EBI did not share specifics on what

adjustments were made but further investigation showed that performance remained stable throughout the year.

In Year 3, we had three ongoing 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). LDM is testing a new version of the file transfer protocol, based on UDP multicast, on a 5-site testbed that includes UCAR, 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 current setup, UCSD and UCAR are having 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 and narrowed the issue down to a single switch in the FRGP network. This device was scheduled to be replaced in July 2020, however, delays in procurement and installation have pushed this installation into 2021.
- **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, 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 will purchase two portable network attached storage appliances (NAS), load them at the Arecibo data center, transport the appliances to the UPR campus, and start transfers using Globus to TACC. While waiting for the appliances to arrive at Arecibo, staff at EPOC, UPR, and TACC performed network baselining with perfSONAR. The testing showed a clean path between the UPR campus and TACC. During Quarter 4 nearly 1TB of data was moved from Arecibo to TACC with ongoing performance engagement and monitoring. 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 are also working with Arecibo staff to create a portal and workflow to allow easy access to the data for researchers.
- **121 - Event Horizon Telescope (EHT):** EPOC staff were pulled into an ongoing conversation between the astronomy researchers that are part of 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 different locations to the MIT Haystack Observatory for consolidation. The telescopes each stream data at a rate of at a rate of 64 Gbps. The NOEMA telescope in France and the 30m telescope in Spain were chosen as the first two facilities to work with and baseline network performance between the sites and MIT Haystack. In Quarter 4, we worked to do baseline testing to MIT Haystack from both telescopes and draw a network diagram that includes traceroute data. This testing included using iPerf, scp, and their current standard data transfer tool, m5copy. EPOC staff also reached out to the staff at the two EHT telescope participants in Hawaii to start baselining their networks as well.

6.B Consultation Cases

In Year 3, we had 64 completed and 17 ongoing Consultations. In Quarter 4, the primary Consultation topics were assistance with grant review responses, Science DMZ architectures, and data transfer strategies.

Completed Consultations included:

- **35 - Kent State University (KSU)** - KSU engineers reached out to EPOC staff for feedback on their draft plans of their upcoming network redesign (including equipment options) and performance issues with their current DTN. A discussion via Zoom and email followed that included examining possible issues with buffering and the need to clarify their data transfer use cases in order to deploy the correct infrastructure.
- **50 - Mississippi State University (MSU)**: MSU staff expressed interest in doing a Deep Dive or an EPOC-lead Deep Dive Training Event. When additional information was requested, MSU reported that they needed to focus on other priorities related to COVID restrictions, and that this case should be closed. They will reach out again when they have time to fully engage.
- **51 - Kent State University (KSU)**: KSU engineers inquired about borrowing the EPOC Viavi network tester in early 2020. During the conversation, the request shifted to using perfSONAR (pS) for testing as the EPOC hardware tester was unavailable. KSU engineers initial attempts with pS were unsuccessful, but EPOC staff was able to provide additional setup documentation along with assistance in troubleshooting the configuration. After several corrections to the configuration, regular testing began working as expected with the exception of throughput testing, which was still failing. Additional investigation by EPOC staff revealed IPv6 communication issues between the pS nodes which were temporarily remedied by forcing all tests to run via IPv4 while campus engineers worked to restore IPv6 connectivity. The end result of this engagement was a completely functional perfSONAR mesh at Kent State.
- **57 - University of Hawaii (UH)**: UH Astronomy staff were seeing poor file transfer performance to and from the Ohio State University (OSU). UH network engineers worked jointly with EPOC staff to troubleshoot the UH DTNs for configuration errors. After installation of a perfSONAR node at OSU that showed consistently good transfer

rates, UH engineers found that the application used by the astronomers for the file transfers used an unpatched version of rsync. Having identified that the file transfer application was the main performance issue, the case with EPOC was closed and UH staff are working directly with the astronomy team to upgrade or adapt the workflow for better performance.

- **62 - Lonestar Education and Research Network (LEARN):** LEARN engineers were asking for help designing a low cost but powerful DTN. They reviewed the information EPOC staff has provided via email and zoom meetings and purchased and deployed 2 DTNs for the LEARN network.
- **67 - Veterans Administration (VA):** Engineers from the VA reached out to EPOC and ESnet staff for additional information related to the basics of SDN. EPOC staff responded, and also engaged staff from the FAUCET SDN and Network to Code teams.
- **69 - University of Texas San Antonio (UTSA):** EPOC and UTSA staff discussed ScienceDMZ architecture, possible bottlenecks, policy, security, and the effect of switch buffer sizes on data transfers.
- **72 - Great Plains Network(GPN):** EPOC staff met with the newly funded GPN CyberTeam awardees and gave an overview of the EPOC project, the Roadside Assistance and Consulting process, and information on a then upcoming University of South Dakota Deep Dive. We expect additional collaborations with this team in the future.
- **73 - Duquesne University (DU):** Engineers from Duquesne and EPOC staff discussed testing methodologies for equipment loaned by vendors for evaluation of the Duquesne ScienceDMZ use cases. The EPOC Viavi network tester was sent to Duquesne so they could stress test and evaluate the new networking hardware for use in their upcoming ScienceDMZ.
- **74 - Allen Institute (AI):** EPOC exchanged emails and had a Zoom call with systems engineers and researchers at the Allen Institute about the basics of ScienceDMZ architecture, network hardware, and file transfer nodes. The Institute engineer is interested in further engagement, however, they are currently addressing other pandemic-related critical tasks.
- **75 - Reed College (RC):** A researcher at Reed was writing a grant to support data transfers for a high-end microscope that produced 4TB datasets. EPOC answered questions on file transfer nodes, Globus, and ScienceDMZ. The proposal for the instrument was submitted and funded (NSF #2103677).
- **77 - Compute Canada (CC):** EPOC and ESnet staff consulted with Compute Canada about the Zeek Intrusion Detection System architecture, switch characteristics, server decisions, and storage requirements.
- **80 - University of Central Florida (UCF):** EPOC met with a researcher at UCF working on COVID-19 testing who wanted to share data more effectively with collaborators. We discussed strategies for data movement, data storage, data sharing issues, and approaches to visualization of collected COVID data.
- **81 - Sun Corridor Network (SCN):** Staff from SCN and major universities including Arizona State, University of Arizona, and Northern Arizona discussed their CC* award (# 1925632) with EPOC staff, and participated in a review of the data architecture, DTN,

perfSONAR, use of intrusion detection systems, and use of firewalls at the major partners. EPOC made recommendations, and the grant participants made changes to some of their data architecture to support the findings. A follow up Case (#83) was generated to specifically address other aspects of the award that pertain to smaller institutions in the region.

- **83 - Arizona State University (ASU):** SCN and collaborators on their CC* award (# 1925632) which includes Arizona State, University of Arizona, and Northern Arizona, discussed ScienceDMZ architecture for small campuses that will be supported via the award. A focus on network security approaches was the overall goal of the proposal, with some of the large facilities aiming to offer services to smaller institutions.
- **84 - South African National Research Network (SANReN):** A SANReN engineer requested help contacting Globus administrators at German Climate Computing Centre (DKRZ), a national climate computing center in Germany. EPOC staff verified that SANReN had the correct contact point for this type of issue and made sure that follow up took place.
- **85 - Maryland Research and Education Network (MDREN):** The CTO of MDREN reached out to discuss a CC* proposal under review (#2018823). A reviewer had requested additional information about the suitability of the switches included in the proposal for large science data flows. MDREN sent a copy of their bill of materials (BOM) and network diagrams to EPOC staff for review. EPOC and MDREN staff discussed the concerns, the proposed hardware, and the proposed architecture.
- **86 - OneNet:** EPOC staff worked with OneNet staff to review questions they received from an NSF review panel about switch buffering, network hardware, and their proposed ScienceDMZ as part of award # 2018453.
- **87 - Pennsylvania State University (PSU):** Staff from PSU reached out to EPOC engineers to get advice on the needed data architecture to support transfers from the Berkeley Advanced Light Source (ALS) for computed tomography (CT) imaging. The consultation included a discussion of expected file transfer rates between PSU and the Berkeley ALS Berkeley. EPOC staff provided baseline perfSONAR tests between the borders of the institutions. The tests showed a very clean and high bandwidth path. EPOC engineers explained, however, that the disk-to-disk transfer rate may be less depending on each institution's network architecture.
- **88 - Oakridge National Laboratory (ORNL):** Engineers from ORNL requested assistance configuring a new perfSONAR MaDDash installation. EPOC staff were able to correct configuration and permission errors. ORNL engineers have expressed interest in further perfSONAR training when it becomes available.
- **89 - Ocean State Higher Education Economic Development and Administrative Network (OSHEAN):** An OSHEAN engineer requested information regarding perfSONAR installation and configuration, along with current best practices for DTN design and implementation. EPOC engineers were able to answer these questions over the course of two separate Zoom meetings.
- **90 - University of California Santa Cruz (UCSC):** Researchers at the Agency for Science, Technology and Research (A*STAR) in Singapore have recently begun transferring large amounts of data from the University of California Santa Cruz (UCSC)

Genome Browser. Slow transfer speeds of less than 60Kbps were noted by IN@IU staff using NetSage during an investigation into possible COVID-19 research data transfer performance issues. Joint investigation by EPOC, UCSC, and A*STAR staff revealed asymmetric routing between UCSC and A*STAR due to a multi-week outage of the I2/SingAREN transpacific circuit. Once the circuit was restored to service, normal routing of traffic between UCSC and A*STAR resumed, and the investigation was halted.

- **91 - Rede Nacional de Ensino e Pesquisa (RNP):** EPOC engineers consulted with RNP, the Brazilian National Research and Education Network (NREN), regarding specifications for 100G capable data transfer nodes and shared known performance results for various configurations.
- **92 - University of Central Florida (UCF):** UCF requested information about sustainability models for both staffing and hardware. A custom spreadsheet of the service and financial models was completed and shared with UCF.
- **93 - Kent State University (KSU):** Continuing efforts which began in 2019 (Cases 35 and 51), a KSU engineer reached out for advice and troubleshooting assistance for their 100G-focused perfSONAR MaDDash deployment. EPOC staff offered suggestions to correct several service-impacting errors in the configuration files, suggested updating the software, and provided additional support via targeted documentation. The KSU pS MaDDash deployment was fully operational on June 3, 2020.
- **96 - Yale:** EPOC and Yale staff met virtually to discuss their proposed ScienceDMZ architecture, DMZ best practices, and known pros and cons of using Cisco SDA network technology. Further discussion centered around Yale use cases and where to best position data transfer nodes.
- **97 - University of Central Florida (UCF):** UCF requested assistance to redesign their research and production networks, including the restoration of their perfSONAR testing infrastructure. Initial contact with EPOC resulted in the development of a basic plan of attack and subsequent email exchanges have helped keep the process moving forward. UCF has decided to step back and make higher level decisions about their network design and implementation. This consultation also resulted in a larger discussion around the EPOC deep dive process and creating science-based use cases for the network changes.
- **98 - 126.com:** A Chinese test engineer emailed with questions about perfSONAR hardware used to generate results on the fasterdata.es.net site. EPOC engineers sent links to the hardware specifications, software used, and configuration tuning.
- **99 - Arizona State University (ASU):** EPOC engineers shared best current operational practices for data transfer node (DTN) design and operation. They also discussed using virtual machines for DTN nodes.
- **100 - European Bioinformatics Institute (EBI):** With the rollout of the EU COVID-19 data portal at EBI, the amount of data shared between EBI and PHS, a consortium of academic hospitals based in Cambridge, MA, increased from 945GB per month in January to over 40TB per month by July. Using NetSage, EPOC staff noticed that these new data transfers averaged below 25Mbps. After discussing the issue and networking set-up with EBI engineers and running tests on the network path between MIT and EBI, we also worked with engineers from the Northern Crossing (NOX) exchange and from

PHS to learn more about the PHS network set-up. From these discussions, engineers at EBI suspected that the bottleneck was likely due to load balancers and/or firewalls being periodically overwhelmed by the increase in traffic. In November, EBI made internal adjustments that helped improve performance from an average of ~25Mbps in Quarter 1 to an average of ~90Mbps in Quarter 3. After multiple attempts to get more information about what specific changes were made by EBI this case was moved to OBE.

- **101 - Lawrence Berkeley National Laboratory (LBNL):** As the COVID-19 pandemic forced LBNL staff to work from remote locations, a number of users started reporting network performance problems reaching lab resources from untested commercial internet connections. EPOC staff worked with an LBNL user to utilize perfSONAR on a laptop via Docker to test their remote network connection's capabilities to LBNL. The user's perfsonar testing via a remote desktop docker container provided baseline throughput measurements to public perfSONAR servers.
- **102 - United States Department of Agriculture (USDA):** USDA and EPOC staff discussed over the course of several zoom sessions and an email thread the Science DMZ security model and how to convince decision makers that it is secure. They would like to provide evidence that the security model for Science DMZ has been proven at other Government entities. Firewalls in the USDA network have been identified as transfer performance bottlenecks, but USDA engineers have not been able to receive approval to use access control lists as an alternative. EPOC engineers provided information on the security of the Science DMZ network architecture, DTNs, and perfSONAR nodes.
- **103 - Rubin Observatory (VRO):** EPOC staff exchanged emails with a Rubin network engineer about using the perfSONAR CLI and directed the engineer to the perfSONAR users group mailing list for more advanced discussion.
- **104 - Allen Institute (AI):** An engineer at the Allen Institute reported performance issues with a science data transfer to a commercial cloud provider via Princeton University. EPOC engineers worked with engineers from both Princeton and the Allen Institute to diagnose the issue. Princeton staff reported in Quarter 3 that they were in the middle of network upgrades and not fully switched over to R&E networking paths, so testing at this time would not provide good results. In addition, staff at the Allen Institute were unable to continue their engagement with the person having the issue, so this Case was closed and will be reopened if the issue continues.
- **106 - Front Range Gigapop (FRGP):** Engineers at FRGP inquired about the importance of larger network memory buffers when using newer versions of the transmission control protocol (TCP) during large data transfers. This included a discussion around these recent protocol developments and how they may reduce buffer importance and overall equipment cost for smaller schools.
- **108 - King Abdullah University of Science and Technology (KAUST):** A KAUST engineer reached out for information on 100G capable data transfer nodes(DTN) and how best to support their researchers. EPOC engineers consulted on how to best tune the new node as KAUST engineers benchmarked disk to disk performance between their new nodes.

- **109 - National Human Genome Research Institute (NHGRI):** Globus connected EPOC staff to a program analyst at NHGRI in the process of setting up a Globus installation and Science DMZ. EPOC staff shared best practices, slide decks, and links via a video conference and email.
- **110 - Front Range Gigapop (FRGP):** FRGP engineers asked for a link to a TCP Bottleneck Bandwidth and Round-trip propagation time (BBR) slide deck referenced during a CI Engineering Lunch & Learn talk. EPOC provided links to various talks on the BBR protocol.
- **111 - Kent State University (KSU):** A doctoral student at KSU had questions about using perfSONAR data for research. EPOC worked with the researcher to understand available data sources as well as the scale and scope of perfSONAR measurement archive data.
- **112 - University of South Carolina (USC):** Staff at USC received a CC* award (#1925484) and were in the process of deploying data transfer and perfSONAR nodes. Engineers from USC connected with EPOC staff to learn about routing and Science DMZ best practices via a video conference and email consultation.
- **114 - New York University (NYU):** EPOC staff were consulting with NYU's Databrary staff about ScienceDMZ's, DTN's, perfSONAR, and using Globus to improve data transfer performance. EPOC shared a few Science DMZ data transfer examples and best practices. Unfortunately, most of the data transfer requests that the Databrary receives were from home connections. We have offered to resume our consultations if they would like to test when they are back on campus.
- **115 - Texas A&M University (TAMU):** TAMU is the lead institution on the NSF funded Building Research Innovation at Community Colleges (BRICCs) award (#2019136). TAMU staff engaged EPOC for a consultation on ScienceDMZ and Deep Dive best practices for community colleges.
- **116 - Colorado State University (CSU):** EPOC staff presented on the Deep Dive process at the NOAA meeting in December highlighting previous work with Cooperative Institute for Research in the Atmosphere (CIRA) at CSU.
- **118 - Northern Arizona University (NAU):** Northern Arizona University, Arizona State University, the University of Arizona, and Sun Corridor Regional Network are collaborators on a CC* award (#1925632) to engage with community colleges on identifying science drivers and providing technical solutions. The proposed ASU Deep Dive was planned in part to address this need but has been delayed due to COVID-19. EPOC staff discussed virtual approaches to the in person Deep Dive as an option for the collaborating groups to pursue. The ticket was closed after the consultation and the parties have not requested further assistance.
- **119 - University of Hawaii (UH):** EPOC staff consulted with University of Hawaii engineers and astronomers on data movement from CERN to UH in support of the researchers working with the Alpha Magnetic Spectrometer on the International Space Station. It was determined that although the researchers at the University of Hawaii has access to well-connected DTNs and high-performance data movement tools including Globus, the particular data set they were interested in receiving from CERN used a specific workflow management tool, the CernVM File System (CVMFS), that operates in

a streaming fashion. The Hawaii team performed steps to set up a local instance of the CVMFS and GridFTP, and requested that the ticket be closed.

- **120 - Yale University (YU):** EPOC and Globus staff met with the Yale Director of Network Architecture to discuss ways to optimize Globus data transfers between DTNs.
- **122 - State University of New York (SUNY):** EPOC and Globus staff consulted with the chief information security officer (CISO) and staff at SUNY to explain the security implications of the Science DMZ model and Globus.
- **123 - Lonestar Education and Research Network (LEARN):** LEARN staff were investigating cyberinfrastructure automation tools that would help with continuous integration and continuous delivery along with configuration management tooling and how the two can work together. They reached out to understand the best current practices in the use of Ansible configuration management and how it can integrate with other automation tool suites.
- **124 - Oregon State University (OSU):** Staff at Oregon State reached out for information about 100Gbps perfSONAR specifications. EPOC staff provided best current practices. EPOC also discussed the need for a 100G capable node versus the use and utility of a 10G node.
- **125 - Wayne State University (WSU):** The office of the CIO from WSU in Michigan, an awardee of several CC* awards (#1925467, #1245719, #1541335), requested EPOC assistance related to the sustainability of CI investments. In particular, they were interested in how legacy infrastructure can be kept up to date without having to pursue external funding on a repeated basis. EPOC staff met with the WSU team to discuss strategies that have been implemented by other institutions and provided a spreadsheet that showed an example service sustainability model. The spreadsheet includes a 5 year life cycle on CI equipment for budget planning, financial numbers for the CI-Engineer staff, and Data Transfer Nodes.
- **126 - King Abdullah University of Science and Technology (KAUST):** A staff member at KAUST asked for help doing a Globus write test to an ESnet host in the US. ESnet staff were contacted and the proper Globus access permissions were set, after which testing was successful.
- **127 - Heinrich-Heine-Universität Düsseldorf (HHU):** An assistant professor at HHU reached out about errors he was receiving while using Globus to test transfers to ESnet DTNs. EPOC staff investigated and found that HHU is not connected to an R&E network and only has commercial connections to the Internet, which caused tests to ESnet DTNs to fail.
- **129 - Oak Ridge National Lab (ORNL):** EPOC staff were contacted by an HPC Network Engineer from ORNL to help resolve lingering issues with their perfSONAR mesh, which consists of two on-campus nodes and two off-campus nodes. The off-campus nodes are located on a military base where additional security protocols have been implemented which make establishing regular testing difficult. Review of the mesh configuration and log files revealed that the necessary ports were being blocked and that the Network Time Protocol (NTP) had not been enabled. Troubleshooting is continuing.
- **130 - Pennsylvania State University (PSU):** The Assistant Director of the Administrative Data Accelerator at PSU attended a talk given by EPOC staff at the

CaRCC Emerging Centers call [34] and afterwards contacted EPOC staff for a consultation on a possible CC* proposal. EPOC provided a letter for the proposal that PSU submitted to the CC* program.

- **131 - Miami University (MU):** After EPOC staff presented at the CaRCC Emerging Centers call [34], a researcher from Miami University requested some clarifying information on the difference between the NSF CC* and MRI programs. A response was sent, with the suggestion to contact the relevant program officers with one-page descriptions as appropriate.
- **132 - Allegheny College (AC):** The Vice President for Information Services at Allegheny College reached out after EPOC staff presented at the CaRCC Emerging Centers call [34] and requested some additional information on the program. EPOC staff shared the links and offered additional help to follow up as needed.
- **136 - Yale University (YU):** The Director of High-Performance Computing at Yale's Center for Research Computing asked for a consultation on configuration and tuning of DTNs as they looked to purchase a new node for their Science DMZ. Discussions in email centered on recommended hardware specifications for the node, basic Science DMZ architecture, and best practices for connecting instruments to storage arrays.
- **137 - Singapore Advanced Research and Education Network (SingAREN):** A technical specialist at SingAREN requested write access to ESnet DTNs to allow them to test a new DTN they had recently installed. EPOC staff put the specialist in touch with ESnet systems engineers to complete the access request.
- **138 - Singapore Advanced Research and Education Network (SingAREN):** SingAREN engineers were setting up 10G DTNs in Indonesia, Thailand, and the Philippines. During initial perfSONAR testing on the DTNs, they noticed that the path to or from ESnet perfSONAR nodes in some cases were taking commercial networks. EPOC staff worked along with ESnet engineers to adjust the routing policy on the ESnet network to enable the SingAREN DTNs to connect and successfully run the perfSONAR tests.
- **139 - Princeton Plasma Physics Laboratory (PPPL):** An engineer at the PPPL asked about best practices for data transfer and perfSONAR nodes. EPOC staff answered questions and shared best practice documents.
- **140 - SingAREN:** A technical specialist at SingAREN contacted EPOC with questions centered around DTN setup and best practices. EPOC engineers shared information from the ESnet fasterdata website and discussed in depth tuning of both disk systems and low level system configurations.
- **144 - Kent State University (KSU):** A KSU engineer inquired about a discoverable list of DTNs similar to the perfSONAR directory. EPOC staff explained that there is no global list of Globus endpoints, however, using the Globus tool a user can search for specific endpoints.
- **149 - Arecibo:** Arecibo staff members are using network attached storage (NAS) devices to be able to download data from the Arecibo Data Center and then drive it to the University of Puerto Rico to take advantage of their R&E network connection back to CONUS, however the performance of the transfers on and off the NAS was very slow. EPOC engineers procured an identical NAS to test and tune the device in the lab and

create a guide for how the Arecibo NAS units should be set up. This tuning provided roughly a ten-fold speedup in data transfer performance.

- **156 - Hong Kong Academic and Research Network (HARNET):** The HARNET network operations center alerted EPOC and TransPAC that a user on the HARNET network could not connect to a network resource in Sri Lanka. Trace routes showed the traffic going via TransPAC, Internet2, the National Knowledge Network (NKN) in India, and the Lanka Education and Research Network (LEARN) which is the Sri Lankan R&E national network. EPOC engineers contacted Sri Lanka LEARN engineers who were able to make changes to how they advertise routes to NKN. These changes were confirmed by the user at HARNET to have fixed the problem.

Ongoing 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. The Deep Dive has been postponed due to COVID-19, however, EPOC staff presented virtually on the process during a NOAA meeting in Quarter 3. EPOC staff are waiting for feedback and scheduling of a discussion with NOAA stakeholders.
- **105 - Lafayette College (LC):** The manager of research and high performance computing at Lafayette College requested information about ScienceDMZs, DTN design, and data architecture best practices. This consultation is ongoing for data transfer testing, network performance testing and Science DMZ design for a potential NSF CC* grant proposal but had slowed due to Lafayette pending projects.
- **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 Quarter 3. NetSage shows that the transfer rate is variable between 500Mbps and 1.3Gbps after the changes. Early in Quarter 4 during a routine check of NetSage it was found that the transfer rates had again fallen under 500Mbps. EPOC and TAMU Engineers discussed the performance drop at that time and TAMU engineers found no apparent changes in the network configuration. TAMU engineers are purchasing and installing a perfSONAR node to continue testing.
- **117 - Allen Institute (AI):** EPOC staff are consulting with a systems engineer at the Allen Institute about currently implemented CI resources in the community and funding opportunities.
- **135 - Texas Advanced Computing Center (TACC):** TACC is in the process of a network refresh and the manager of Network Operations has asked for EPOC consultation during the planning stages.
- **142 - University of Alaska Fairbanks (UAF):** EPOC staff consulted with engineers from the International Arctic Research Center at UAF on applying for a CC* grant. Emphasis was placed on science use cases, which area of the solicitation to apply to, Science DMZ best practices, data transfer nodes, and the need for deep buffered switches to support long distance data transfers. UAF has challenges in establishing

connectivity to the lower 48 states and is weighing connectivity options for their WAN connection versus the local needs on the LAN. The engagement continues as UAF engineers prepare their grant and are targeting the October 2021 deadline.

- **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. The facility is still seeking funding but is currently evaluating the proper setup of a data architecture: network infrastructure, storage, computational power to calibrate and analyze samples, and the workflow tools to be used to control the components.
- **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 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, which provides the network connection to NRAO.

EPOC staff are also helping the group evaluate new transfer tools like Globus that may speed up the overall data transfer. This consultation is ongoing.

- **152 - Brown University (BU):** A senior network engineer at Brown University asked questions about adding a firewall to the path of an established Science DMZ. EPOC staff are answering questions and sharing best security practices via email.
- **153 - Saint Louis University (SLU):** EPOC staff met with the director of network services at SLU to discuss support and services that the EPOC project can provide. SLU staff are interested in more information on Deep Dives, how to identify science use cases on campus, and the Data Movement Exhibition. This engagement will continue as SLU staff reach out internally to start identifying use cases.
- **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 are currently discussing NetSage 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 to address a number of campus use cases. 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:** SU runs a DDoS scrubbing service at the border of their campus network that requires special configurations for packet size manipulation and redirecting traffic to the scrubbing services network. The Chief Network Architect at NYSERNet reached out to EPOC for help investigating the performance impact this has on traffic on the campus network, as many of their connector universities are also signing up for this type of scrubbing service. EPOC engineers are working with the NYSERNet staff to test the effect this service may have on data transfers using perfSONAR.
- **158 - New York University Langone Health (NYULH):** Staff at NYULH are investigating the ScienceDMZ model and how it might fit their network. A meeting has been scheduled with EPOC staff in Year 4 Quarter 1 to discuss best practices, architecture, and DMZ security practices.
- **159 - Yale:** The program director for the next generation network project at Yale reached out to discuss some performance abnormalities they are experiencing with Globus file transfers. EPOC engineers have started initial troubleshooting and will continue this engagement into the next quarter.
- **160 - American Museum of Natural History (AMNH):** EPOC staff did a NetSage demonstration for AMNH and found example transfers with very poor performance including a 2.6TB transfer to the University of Arizona that took over a week with average performance of 4.6 Mbps. EPOC engineers will start troubleshooting this performance issue in the next quarter.
- **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 specifically at RPI to support the use of a commodity network protection service. The concern was that network hardware manipulating packet sizes independently of end devices, to conform to the specified MTU configuration of an external service, would impact the performance of data transfers by inducing packet fragmentation, at best, or MTU blackholes, at worst. Investigation into these concerns were not conclusive due to the 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 are able to handle the MTU changes with limited impact to overall data transfer speeds. EPOC engineers continue to work with NYSERNet staff to test these MTU issues against a NYSERNet perfSONAR node.

When a Roadside Assistance Case or Consultation does not respond to multiple attempts to move it forward over a more than six-week time frame, we consider those Cases Overtaken by Events (OBE). In Year 3 we had 1 OBE Roadside Assistance Cases.

- **94 - National Library of Medicine (NLM):** Researchers at ASGC in Taiwan have been consistently transferring roughly 3TB of biology-research data per month from the National Library of Medicine (NLM). While adequate performance was noted on some

flows, the vast majority showed poor performance with transfer speeds of less than 5Mbps. EPOC staff engaged with engineers from NLM and ASGC. The first challenge was to determine the route for the data flows, as the Academia Sinica network is subdivided into segments that reach the United States via different paths depending on their source location in the network. ASGC engineers were able to determine that the endpoint of these transfers is the Agricultural Biotechnology Research Center in Southern Taiwan, and further analysis suggested that routing changes internal to the Academia Sinica network may yield improved performance. This case started on June 18, 2020, and moved to OBE status on December 9, 2020, after EPOC staff were unable to re-engage after multiple escalation attempts.

6.C Additional Outcomes and Plans for Year 4

In Year 3, the top topics included:

- Transfer Performance: 21 Tickets
- Science DMZ: 19 tickets
- DTN: 19
- CC* Proposal / Award related: 14

Based on the broad set of Roadside Assistance and Consultation Cases addressed during Year 3, we also created the following Best Practice documents:

- In response to Cases #62, 91, and 108, a document to help engineers evaluate the advantages and disadvantages of deploying a 100G DTN <https://epoc.global/wp-content/uploads/2020/12/2020-12-07-I-think-i-need-a-100G-perfSONAR-node-2.pdf>.
- In response to Case 124 and discussions on the CI Engineering List, a document to help engineers evaluate the advantages and disadvantages of deploying a 100G perfSONAR node, <https://epoc.global/wp-content/uploads/2020/12/I-think-I-need-a-100G-DTN.2020-12-4.pdf>
- In response to Cases #91, 99, 105, 108, and 136, we are also working on a training video and accompanying documentation, for how to select a 10G DTN (see Section 11).
- In response to Cases #1, 23, 94, and 138, we are working on BGP Best Practice documents, listed in Section 11.B.

Over the course of Year 3, we also have eighteen presentations that walked through components of the Roadside Assistance Process and talked about the results. We envision this continuing in Year 4, and also plan to expand the documentation to include:

- 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.
- Several two-page documents aimed at engineers that discuss the more common issues EPOC staff see during a Roadside Assistance Case, such as MTU/packet size issues seen during data transfers and how to identify common causes of packet loss on a campus network.
- A guide to understanding why a 100G connected perfSONAR or DTN node may not always achieve 100G transfer rates.

During Year 4, we plan to pivot to include a more proactive approach to Roadside Assistance and Consulting as well, via several different methods. These include:

- We plan to contact each CC* awardee to introduce EPOC and more specifically offer our consultation services.
- The Data Mobility Exhibition (DME) also continues to run in 2021 and we will reach out to each participant to discuss their test results and offer the roadside assistance program as a method to improve their data transfer results.
- Working with the TransPAC, NEA3R and NetSage projects we will continue to identify poorly performing international data transfers and engage with the end users and networks along the path to increase the performance.

We also plan to be a larger part of the XSEDE/ACCESS and offer to help with any issues seen transferring data between the sites.

6.D Metrics

Table 1: A summary of Year 3 and Quarter 4 Roadside Assistance and Consultation Cases. 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)	Size: S, M, L
35	KSU	OARnet	Cons	9/24/19	8/20/20	DTN	Eng	Infra	R	S
50	MSU	GPN	Cons	11/1/19	6/1/20	DD, Training	Eng	Infra	E	L
51	KSU	OARnet	Cons	11/1/19	12/2/20	Tester, PS	Eng	Infra	R	S
57	UH	HI Astro, OARnet	Cons	12/6/19	5/26/20	Trans Perf	Sci	Infra	E	L
59	SLU	GPN	RA	12/11/19	12/1/20	Trans Perf	Sci	Bio	E	S
62	LEARN	LEARN	Cons	12/13/19	1/5/21	DTN	Eng	Infra	O	-
63	NOAA	FRGP	Cons	12/16/19		DD	Eng	Infra	O	-
67	VA	-	Cons	1/9/20	5/11/20	SDN	Eng	Infra	O	-
69	UTSA	LEARN	Cons	1/21/20	8/20/20	DMZ, Security	Eng	Infra	R	L
71	TTU	LEARN	RA	2/6/20	6/23/20	Trans Perf	Eng	Climate	E	L
72	GPN	GPN	Cons	2/13/20	6/22/20	Intro, DD	Eng	Infra	E	L
73	DU	KINBER	Cons	2/20/20	10/19/20	Tester	Eng	Infra	E	S
74	AI	CENIC	Cons	2/28/20	1/20/21	DMZ, DTN	Eng	Bio	R	S/M
75	Reed	Globus	Cons	3/5/20	3/31/21	DMZ, DTN, Globus	Eng	Bio	E	S
76	NCAR	FRGP	RA	3/6/20		Trans Perf	Eng	Climate	R	L
77	CC	-	Cons	3/13/20	5/13/20	IDS	Eng	Infra	R	L

80	UCF	-	Cons	3/31/20	5/11/20	Arch	Eng	Bio	R	L
81	SCN	SCN	Cons	4/3/20	11/18/20	DMZ, PS, Grant, DTN, IDS	Eng	Infra	E	L
83	ASU	SCN	Cons	4/6/20	8/20/20	DMZ, Security	Eng	Infra	E	S
84	SANReN	NEAAR	Cons	4/3/20	4/8/20	Trans Perf	Eng	Infra	R	M
85	MDREN	-	Cons	4/23/20	6/30/20	Trans Perf Grant, DMZ	Eng	Infra	E	
86	OneNet	GPN	Cons	4/29/20	9/3/20	DMZ	Eng	Infra	E	
87	PSU	KINBER	Cons	4/30/20	9/30/20	Trans Perf	Sci	Bio	E	L
88	ORNL	SoX	RA	5/5/20	5/29/20	PS	Eng	Infra	R	L
89	OSHEAN	-	Cons	5/11/20	6/10/20	PS	Eng		R	S
90	UCSC	CENIC, TP	RA	5/20/20	9/16/20	Trans Perf	Eng	Infra	R	L
91	RNP	-	Cons	5/27/20	8/20/20	Trans Perf	Eng	Infra	R	L
92	UCF	-	Cons	5/29/20	10/28/20	Sustain	Eng	Infra	E	L
93	KSU	OARnet	Cons	6/1/20	6/03/20	PS	Eng	Infra	E	S
94	NLM	TP	RA	6/5/20	12/9/20	Trans Perf	Eng	Bio	R	S
95	CCNY	TP	Cons	6/8/20	3/17/21	Trans Perf	Eng			
96	Yale	-	Cons	6/16/20	9/28/20	Arch, DMZ	Eng	Infra	E	L
97	UCF	-	Cons	6/26/20	9/25/20	PS	Eng	Infra		S
98	126.com	-	Cons	7/13/20	7/14/20	Perf	Eng	Infra	O	
99	ASU	SCN	Cons	7/13/20	8/26/20	DTN	Eng	Infra	R	L
100	EBI	NEAAR	RA	7/16/20	10/21/20	Trans Perf	Eng	Bio	R	L
101	LBNL	CENIC	RA	7/17/20	2/22/21	PS	Eng	Infra	R	L
102	USDA	-	Cons	7/24/20	2/22/21	DMZ, Security	Eng	Infra	R	L
103	VRO	-		8/3/20	9/16/20	PS	Sci	Astro	O	S
104	AI	CENIC	Cons	8/3/20	1/20/21	Perf	Eng	Infra / Bio	R	L
105	LC	KINBER	Cons	8/13/20		DMZ, DTN	Eng	Infra	E	S
106	FRGP	FRGP	Cons	8/21/20	9/28	DMZ	Eng	Infra	E	S/M
107	Arecibo	TACC	RA	8/21/20		Trans Perf DMZ	Eng	Astro	R	L
108	KAUST	NEAAR	Cons	8/26/20	2/1/21	DTN	Eng	Infra	R	L
109	NHGRI	Globus	Cons	8/31/20	12/2/20	DMZ, Globus	Eng	Bio	R	L
110	FRGP	FRGP	Cons	9/2/20	9/8/20	Perf	Eng	Infra	R	L
111	KSU	OARnet	Cons	9/8/20	12/9/20	pS	Sci	CS	R / E	S / M
112	USC	-	Cons	9/10/20	1/5/21	DTN				

113	TAMU	TACC, LEARN	Cons	9/10/20		Perf	Eng	Infra	R	L/M
114	NYU	Globus	Cons	10/5/20		DMZ, DTN, PS, Globus	Sci	Infra	R / E	
115	TAMU	LEARN	Cons	10/15/20	2/1/21	DMZ, DD	Eng	Infra	E	S/M
116	CSU	FRGP	Cons	10/20/20	11/30/20	DD	Sci	Climate	R	M
117	AI	CENIC	Cons	10/21/20		Research	Eng	Infra	R	L
118	NAU	SCN	Cons	10/22/20	2/8/21	DD	Eng	Infra	E	S
119	UH	U of Hawaii Astro	Cons	10/22/20	2/1/21	Trans Perf	Sci	Astro	R	L
120	YU	Globus	Cons	10/27/20	2/1/21	Globus	Eng	Infra	R / E	M
121	EHT	NEAAR/TP	RA	10/28/20		Trans Perf	Sci	Astro	R	L
122	SUNY	Globus	Cons	10/28/20	2/22/21	DMZ, Globus	Eng	Medical	R	L
123	LEARN	LEARN	Cons	10/28/20	2/3/21	Automation	Eng	Infra	R / E	L
124	OSU	-	Cons	11/2/20	2/1/21	PS	Eng	Infra	R / E	L
125	WSU	-	Cons	11/2/20	2/1/21	Sustain	Eng	Infra	R / E	M
126	KAUST	NEAAR, Globus	Cons	11/9/20	12/7/20	DTN, DME, Globus	Eng	Infra	R / E	L
127	HHU	NEAAR, Globus	Cons	11/13/20	12/8/20	Globus	Eng	Infra	R / E	S
128	EBI	NEAAR	RA	11/5/20	1/13/21	Trans Perf	End	Bio	R	
129	ORNL	SoX	Cons	12/7/20	1/27/21	pS	Eng	Infra	R	M
130	PSU	KINBER	Cons	12/17/20	3/1/21	Grant	Eng	Infra	R	L
131	Miami	OARNet	Cons	12/17/20	12/17/20	Grant	Eng	Infra	E	S
132	AC	KINBER	Cons	12/17/20	12/17/20	Grant	Eng	Infra	E	S
133	UA	GPN	Cons	12/18/20	2/9/21	Grant	Eng	Infra	R / E	S / M
134	UMiss	-	Cons	12/21/20	2/9/21	Grant	Eng / Sci	Bio / Healthcare	R	S / M
135	TACC	TACC	Cons	1/5/21		Arch	Other	Infra	E	L
136	Yale	-	Cons	1/11/21	2/8/21	DTN	Eng	Infra	R/E	M/L
137	SingAREN	TP	Cons	1/14/21	1/26/21	DTN	Eng	Infra	R	M/L
138	SingAREN	TP	Cons	1/14/21	1/19/21	DTN, Routing	Eng	Infra	R	M/L
139	PPPL	-	Cons	1/15/21	1/21/21	DTN, pS	Eng	Infra	R	M / L
140	SingAREN	TP	Cons	1/19/21	2/4/21	DTN	Eng	Infra	R / E	M / L
141	ASU	SCN	Cons	1/22/21	2/17/21	Grant	Eng	Infra	R / E	L
142	UAF	CENIC	Cons	1/28/21		Grant	Eng	Infra	R / E	M
143	ASU	SCN	Cons	2/9/21		Arch, DD	Eng	Infra / Bio	R / E	L

144	KSU	OARnet	Cons	2/9/21	2/9/21	DTN	Sci	Infra / CS	R / E	S
145	LEARN	LEARN	Cons	2/11/21	2/11/21	Grant	Eng	Infra	R / E	L
146	USC	CENIC	Cons	2/11/21	2/18/21	Grant	Eng	Infra	R / E	L
147	SARAO	-	Cons	2/16/21		Trans Perf	Eng	Infra	R	L
148	OneNet	GPN	Cons	2/16/21	2/16/21	Grant	End	Infra	R	
149	Arecibo	-	Cons	1/29/21		Trans Perf	Eng	Astro	R	
150	UNC-C	-	Cons	2/18/21	2/19/21	Grant	Eng	Infra	E	S
151	OARnet	OARnet	Cons	2/19/21	2/19/21	Grant	Eng	Infra	R	L
152	BU	-	Cons	2/24/21		DMZ, Firewall, DTN	Eng	Infra	R / E	M
153	SLU	GPN	Cons	2/24/21		Intro, DD	Eng	Infra	E	M
154	NYSERNet	-	Cons	3/4/21		NS	Eng	Infra	R	
155	TAMUSA	LEARN	Cons	3/8/21		DMZ	Eng	Infra	E	S
156	HARNET	TP	Cons	3/9/21	3/19/21	Trans Perf	Eng			
157	Syracuse	-	Cons	3/25/21		Routing, Security	Eng	Infra	E	
158	NYULH	-	Cons	3/26/21		DMZ	Eng	Medical	R	M / L
159	Yale	Globus	Cons	3/31/21		Trans Perf, Globus	Eng	Infra	R / E	L
160	AMNH	SCN	Cons	3/31/21		Trans Perf	Eng	Infra	0/E	-
161	RPI	-	Cons	3/31/21		Trans Perf	Eng	infra	R / E	M

7. 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/2019/04/Application-Deep-Dive-Description-1.pdf>. Jason Zurawski is the lead for this area.

7.A Shift to Virtual Deep Dives

EPOC engaged in a number of conversations internally, with community members, and with our advisory board, to evaluate strategies to adjust our approach to Deep Dives to work in a virtual environment due to pandemic-related travel restrictions. The new approach involves a video session to “train the trainers” and help the local IT staff understand the Deep Dive structured conversation approach so they can work with the individual science groups to fill out the application Case Studies. Once the Case Studies are collected, we schedule a series of Focus

Group video calls, each no more than 2 hours in length, to walk through a subset of the Case Studies and try to identify the CI needs and requirements. The full set of participants then update the Case Studies, and we combine that data and observations from the focus groups into a report that is reviewed and then made public. In Year 3, we held virtual Deep Dives with the University of Central Florida (UCF), University of South Dakota (USD), NOAA, and Texas A&M University San Antonio (TAMUSA).

At times, the Virtual Deep Dives have advantages over the in-person version in terms of flexibility. An example of this was with an archeology and anthropology researcher from UCF who typically spends significant time in a remote international location conducting field work. They were able to have the time to think critically about prior workflow and determined that it could be more effective if it used advanced IT support. The old process included hand carrying an entire summer's worth of data back to the lab using removable media, which was then processed over a number of months. After talking with EPOC and the sponsors at UCF, the team is considering ways to use remote data transfers and integration with local computational resources, which in turn could facilitate a parallel workflow during the time of remote study and enable more flexible field work through the availability of partial results.

Once travel restrictions are lifted, we will engage with other sites to see if we can return to hybrid events that feature an in-person component. For now, these approaches are still a work in progress, and are expected to continue.

7.B Completed Application Deep Dives

The following Deep Dive activities are complete and published in Year 3:

- **University of Wisconsin:** The University of Wisconsin requested a campus-wide Deep Dive to assist campus leadership in understanding upcoming CI needs by researchers in high energy physics, space sciences (including support for several NASA and NOAA missions), polar studies (including the IceCube project), bioinformatics, and high throughput computing. The final report was published in May, 2020 [7].
- **LEARN and Baylor University:** In June 2019, EPOC began a conversation with Baylor University about a campus-wide Deep Dive to be jointly run with the LEARN regional network. This event occurred January 6-7, 2020, in Waco, TX. The findings of this report focus on a growing number of data-centric use cases, all of which are heavy users of campus and regional HPC/HTC resources, will help to justify future networking requirements. The final report was published in January 2021 [39].

7.C In Progress Application Deep Dives

The following Deep Dives activities are in progress:

- **University of Central Florida (UCF):** 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 is to produce content that can be used to justify a CC*

proposal for campus in October 2021. UCF is still interested in hosting an eventual live event in late 2021.

- **University of South Dakota (USD):** Staff from UCD 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. The goal is to produce content that can be used to justify a CC* proposal for the region in October 2021.
- **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 started a Virtual Deep Dive to gather scientific use cases and will complete this activity in summer of 2021.
- **NOAA Marine Mammal Acoustics (NOAA NMFS):** 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. The effort will start in spring of 2021 and extend through summer of 2021.

7.D Related Activities

With nine completed Deep Dives to date (and four in progress), we now have a set of data regarding CI preparedness to help us scope future activities. This will include documenting:

- Best common practices for how institutions support specific kinds of research;
- Lessons learned, both positive and negative, for research and technology intersections;
- Common technology gaps;
- Emerging trends for scientific and research use cases.

The first discussion of these points was held at the Quilt Winter Meeting in February, 2020. The roundtable discussion helped to motivate potential services that the R&E networking community had been considering. This documentation, when complete, will be published and shared with the greater R&E community to better influence the design and support strategy that technology can offer to R&E use cases. Eight presentations were given on the topic of Deep Dives, including ones for CARCC [17] and the Quilt [22], and NOAA N-Wave [28] (this is detailed in Consultation case #116) on how the Deep Dive process works.

7.E Upcoming Deep Dives and Year 4 Plans

Deep Dive planning typically involves a series of meetings and conversations over several months with the target institutional leadership and research community. After the event, the EPOC team, joint with the participants, produces a report of the events that can be used by the campus and/or regional network to influence future directions of technology support.

Prior to COVID-related travel restrictions, there were two Deep Dives in planning:

- **Arizona State University / Sun Corridor Network:** In August, 2019, Arizona State University reached out to EPOC to host a potential Deep Dive of campus and regional

requirements. Discussions at the end of 2020 have indicated there is still a desire to reschedule this for 2021, but no firm date has been decided.

- **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 stalled due to COVID-19.

Deep Dive planning typically involves a series of meetings and conversations over several months with the target institutional leadership and research community leading up to an in-person event with all involved parties. After the event, the EPOC team, joint with the participants, produces a report of the events that can be used by the campus and/or regional network to influence future directions of technology support.

Once travel restrictions are lifted, we will engage with other sites to see if we can return to hybrid events that feature an in-person component. For now, these approaches are still a work in progress, and are expected to continue.

7.E Metrics

*Table 2: Metrics for Deep Dive activities in Year 3. All * dates are tentative.*

Meet Date	Appl name	Public/Private	Audience	Offered or Req	Head Count	Issues Identified	Complete Date
June 2019	University of Wisconsin	Priv	University of Wisconsin staff and faculty	Req	35	Instrument Workflows, storage, network upgrades	May 2020
1/6-7/20	Learn, Baylor - 7 Use Cases	Priv	LEARN staff & Baylor faculty and staff	Req	25	Campus capacity upgrades, storage, wide-area data transfer assistance	Jan 2021
Early 2022	Arizona State Univ, Sun Corridor	Priv	ASU, UofAZ, NAU, and Sun Corridor Network staff	Req		TBD	TBD
Virtual - 2020/2021	Univ Central Florida	Priv	University Researchers & Staff, Florida Lambda Rail	Req	20	Increased HPC, access to network testbeds, storage.	Expected Fall 2021
Virtual - 2020/2021	Univ South Dakota	Priv	Staff from GPN, USD, SDSU, Black Hills State, and other guests	Req	40	Regional compute and storage, CI expertise, software integration	Expected Summer 2021
Virtual - 2021	TAMUSA	Priv	University staff and faculty	Req		TBD	Expected Winter 2022
Virtual - 2021	NOAA N-Wave	Priv	NOAA N-Wave staff and scientists	Req		TBD	Expected Fall 2021
2022	Oregon State	Priv	OSU, UofO, and Link Oregon Staff	Req		TBD	TBD

8. 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/2019/09/Network-Analysis-2-pager.pdf>. Jennifer Schopf and Dan Doyle jointly lead this activity.

8.A NetSage Development

Different components of NetSage can be deployed in different ways, depending on the requirements of the customer. During Year 3 of EPOC, the NetSage development team released versions 1.2.0 through 1.7.0, each of which 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 1.2.0 included better linking between Dashboards to ease navigation as well as a general cleanup and unifying of the look and feel of the Dashboards.
- NetSage 1.3.0 brought a new Advanced Flow Analysis Dashboard targeted at advanced users with a greater control over selection criteria, as well as the ability to associate SNMP and flow data.
- NetSage 1.4.0 introduced a new map visualization designed to show where science transfers were happening across the globe.
- NetSage 1.5.0 added a Dashboard focused on science projects to more easily focus on any organizations involved in the project. It also included a Top Talkers over Time Dashboard designed to allow users to understand how the top organizations evolve over longer periods of time.
- NetSage 1.6.0 added the ability to dive into an individual flow and see all of the statistics about that flow. It also added the ability to omit known test traffic, such as from perfSONAR data, to let users look at non-test traffic more easily.
- NetSage 1.7.0 added more tools on the Advanced Flow Analysis Dashboard to allow for discovery of underperforming flows by letting users filter on both overall flow size as well as flow transfer rate.

8.B Current Deployments

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

- **CENIC/PacificWave:** The CENIC/PacificWave 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/>.
- **Front Range GigaPop (FRGP):** The FRGP flow data deployment of NetSage was made public in December 2019 at <https://frgp.netsage.global/>.
- **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.
- **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>. At the end of Year 3 Quarter 3, I-Light provided a list of netblocks in their AS, which they have assigned to their member organizations, based on their Shared Whois Project (SWIP) data. We began using this list to change the organization names of sources and destinations using IPs in those blocks to those of the members to better identify and understand traffic with organizations that do not have their own IP space.
- **KINBER:** Collection of flow data for the PennREN network began at the end of October 2019 and remains publicly accessible at <https://pennren.netsage.global/>.
- **LEARN:** At their 2019 All Hands Meeting, LEARN staff expressed an interest in moving forward to deploy NetSage for the state of Texas network. We made contact with the new LEARN CEO, who has delayed this meeting until later into 2021.
- **OARnet:** We met with OARnet's new CEO in January. 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.
- **Southern Crossroads:** 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/>. We presented at the SoX All Hands Meeting in October with good responsiveness from the audience. We continue work with SoX to expand the datasets that we are collecting to include SNMP.
- **Sun Corridor Network:** An initial NetSage deployment for Sun Corridor was completed in December 2020. This deployment required some additional technical followups to correct for sampling rates which was also completed. We met with the Sun Corridor technical committee to present their data in March 2021, after which we made this publicly available at <https://suncorridor.netsage.global/>. They expressed additional interest in using SWIP and adding SNMP data.
- **TACC:** TACC Flow data has been available since July 2020 and is accessible at <https://tacc.netsage.global/>. We continue to engage with the TACC engineers as needed to ensure they upgrade to the latest release of the pipeline code.

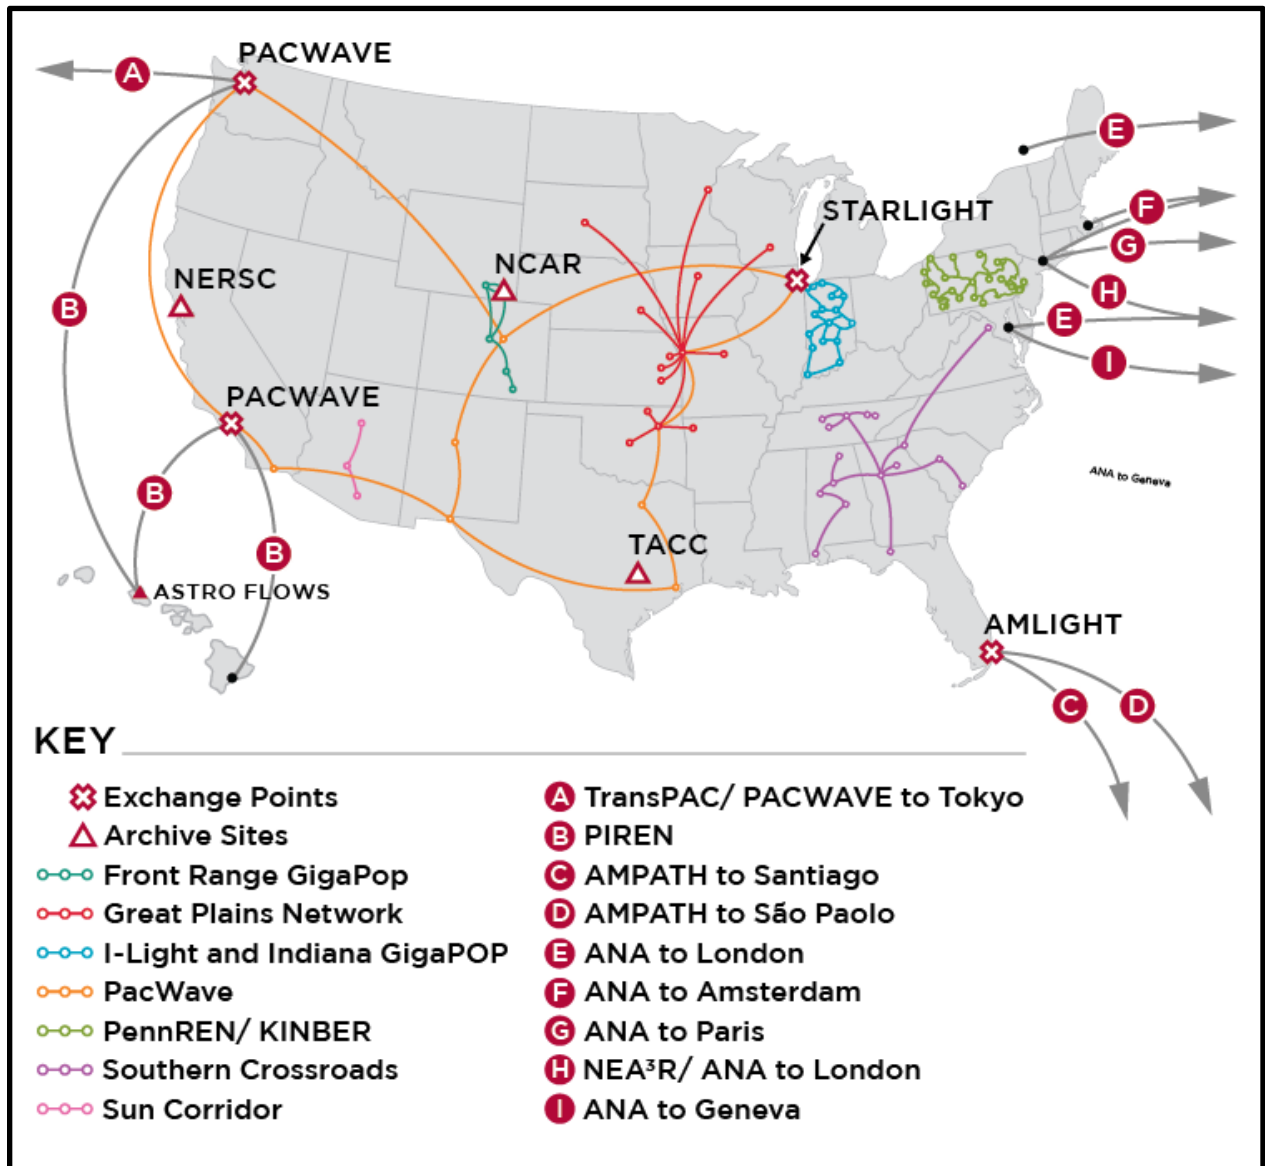


Figure 3: NetSage deployments as of April 2020.

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. No major changes have been made on the EPOC side and this work is running in a stable state.
- **University of Hawai'i Astronomy:** This installation is running in a stable state.
- **NCAR/FRGP:** A Tstat archive has been running since July 2019.
- **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.

These are also shown in the map in Figure 3.

8.C Network Performance Detection

In Year 3, EPOC staff used NetSage to actively investigate performance issues for data transfers related to COVID-19 research. This is an extension of the original project milestone of using NetSage to detect or analyze network “disturbances”. In the many large flows we’ve looked at, we’ve found that increases in COVID-19 data volume often correspond with decreases in transfer performance. Investigations we pursued in Year 3 are described as part of Roadside Assistance Cases #94, 100, and 128.

8.D Additional Outcomes and Plans for Year 4

During Year 3, we gave twelve talks that had a NetSage-centric focus. During Year 4, we plan to continue that type of outreach, to continue to support the domestic deployments of NetSage, and hope to add deployments for new partners. Of special interest is PIREN, which was previously an IRNC partner.

EPOC staff will continue to use NetSage to actively identify and investigate performance issues on large data flows. While the initial work will continue to focus on flows related to COVID-19 research, we expect these efforts to expand to include other science domains where large increases in data volume are detected.

8.E Metrics

Table 3: Metrics for NetSage activities for Year 3.

Where Regional	Data	Date Live	# Monitored Devices	# Large Flows	# Unique Src Orgs	# Unique Dest Orgs
CENIC	SNMP, Flow	2/18	16 routers	364,182,694	4,400	8,612
GPN	SNMP, Flow	10/18	2 routers	78,311,039	3,997	4,797
iLight	Flow	4/19	5 routers	231,854,71	5,860	16,158
KINBER	Flow	11/19	2 routers	72,659,396	4,054	6,324
FRGP	Flow	1/20	1 router	235,154,808	4,908	8,120
SoX	Flow	7/20	3 routers	80,061,207	3,394	2,973
TACC (LEARN)	Flow	7/20	1 router, 4 head nodes	4,041,777	1,130	1,044
TACC (LEARN)	Tstat	1/19	6 head nodes	3,417,044	138	167
UHawaii Astro	Tstat	5/19	1 DTN	1,268,550	207	697
NCAR (FRGP)	Tstat	7/19	1 DTN	70,474,070	642	2,838
NERSC	Tstat	3/18	10 head nodes	5,975,956	171	206

9. Data Mobility Exhibition

The Data Mobility Exhibition provides institutions with a neutral, third-party set up to evaluate a measurable baseline of data transfer performance. Many institutions have been awarded NSF CC* 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 of data in an hour**, which is equivalent to 2.5 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 set up consists of three remote test sites as well as two cloud sites. These well-tuned sites host datasets that range in number from a single file to over 100,000 files, and the file sizes vary from 100MB to 5TB. When the users have completed moving the datasets to and from their institutions, they receive a summary of the upload or download speeds in MB/s rates. EPOC has been pushing the guideline that institutions should be able to move a Terabyte in an Hour, which works out to 277 MB/s.

In Quarter 1, we decided to formally add an Activity area to group the work on the Data Mobility Exhibition. 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/>.

9.A Year 3 Activities

During Year 3, we gave five presentations that included details of the Data Mobility Exhibition. We expect this to increase in Year 4 as we increase the information around moving a Terabyte in an hour. For the DME, 4,154 tests were performed between April 2020 and March 2021. These included tests by member institutions of our Regional Networking Partners, including CENIC, FRGP, GPN, KINBER, LEARN, SoX, and Sun Corridor.

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. From the testing data, we are now advocating that every institution should be able to transfer a TeraByte of data in an hour, and are offering assistance to those who cannot. During the Year 3 tests, the average transfer speed was 0.896 TB/hr, which is also 1.991 Gb/s or 248.89 MB/s.

On the technical side, three of the eight DME testing endpoint sites currently have limited access to the test files, but we are working with the Globus team to reestablish those endpoints. In addition, a set of DME automation scripts were developed to assist with this 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.

9.A Year 4 Plans

In Year 4, we plan to increase engagement through data transfer testing procedures with troubleshooting guidelines/assistance, sharing the current guideline of transferring a Terabyte in an hour disk-to-disk as the baseline for sites as a standard performance expectation, and to showcase clustered/multiple DTNs performance better than single, larger DTNs.

10. 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.

EPOC originally targeted four examples of Managed Services:

1. **perfSONAR**: a widely deployed test and measurement infrastructure that is used by science networks and facilities around the world to monitor and ensure network performance.
2. **Science DMZ**: Dedicated network infrastructure specifically configured for the security and performance implications required for scientific use cases.
3. **Data Transfer Hardware & Software**: PC-based Linux servers built with high-quality components and configured specifically for wide area data transfers along with software layers that can facilitate easier forms of data sharing
4. **Network Testset**: Specialized hardware used to provision and validate network infrastructure.

More information about the Managed Service activity is available online at <https://epoc.global/wp-content/uploads/2019/09/Managed-Services-2-pager.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.

10.A Current Status of Managed Service Deployments

All current engagements are on hold. These include:

- **FRGP:** With their collaborators in the Tribal College consortium and WestNet, FRGP staff members are evaluating the perfSONAR Managed Service. In 2019, EPOC sent six small perfSONAR nodes to be used for a measurement and monitoring deployment for the Tribal Colleges who are in a joint project with FRGP, which are now part of a MaDDash available online. EPOC has been asked to assist with a future training activity, but it is possible that a more effective approach would be to work with FRGP to run the perfSONAR nodes as a Managed Service deployment on behalf of the Tribal Colleges, especially as the deployment expands.
- **LEARN:** LEARN received a CC* award (NSF#1925553) in 2019 to explore the installation of several managed services. As such, they have worked with EPOC to investigate the installation of DMZ, perfSONAR, and DTN hardware at 5 pilot sites.

The one Managed Service activity that has continued over the last project year is our work to deploy perfSONAR nodes jointly with KINBER at Arcadia University. However, 8 months ago, the KINBER organization changed its business model, so our approach shifted to enable Arcadia to be more self-sufficient. During Quarter 3, EPOC staff worked with Arcadia to establish a testing methodology, identify viable test resources in the region, and configure appropriate firewall and ACL rules to allow testing on the Arcadia R&E network. This work concluded in Quarter 4 as initial tests to partners in the region were successful and no further assistance from EPOC was requested by Arcadia.

10.B Service Development - Modern Research Data Portal

The Modern Research Data Portal (MRDP), as detailed at <http://es.net/science-engagement/technical-and-consulting-services/modern-research-data-portals/>, is a design pattern that makes use of the Science DMZ model and DTNs to scale up the data transfer functionality of a data portal. When the data portal gives the user references to data objects, the references point to a well-configured DTN (or DTN cluster) in a Science DMZ, typically using a data transfer platform that can perform job management, fault recovery, and other modern functions.

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 Endpoints.

Initial investigation of this work started in early 2020 with a deployment at ESnet. The initial trials of the Pilot Portal were with the University of Hawaii Astronomy group and LEARN member Baylor. Early results have focused on technical implementation details (e.g. ease of deployment

and integration with existing campus infrastructure). Extensive testing with scientific users will start in Year 4. Initial tests with the Pilot have been challenged due to the complicated install and set up procedure, and work with end users is on hold while we work with members of the Globus team to simplify our approach. Globus has containerized the Globus Connect Server, and we plan to reach out to the pilot sites to gauge their interest in continuing this work. Globus is still in the process of updating the Modern Research Data Portal.

10.C Plans for Year 4

EPOC will be re-evaluating managed services with partners in 2021 and 2022, to gauge the level of participation in deployment and operation. The pandemic has changed priorities for several of the EPOC partners, and further discussion will help EPOC evaluate interest.

11. Training

EPOC is continuing the successful training that ESnet and IU lead 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 Application Deep Dives. All EPOC training materials are available online, including lecture materials, exercises, and recorded sessions when possible.

11.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 of this work 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, which 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, which are meant to be completed by attendees at their own pace after a workshop ends.

EPOC staff participated in the Training Workshop for Educators and Network Engineers on High Speed Network Protocols and Security, (http://ce.sc.edu/cyberinfra/workshop_2020.html) that took place on May 4, 2020. Talks included:

- Zurawski, J., "Science DMZ and TCP",
- Chevalier, S., Southworth, D., "perfSONAR Overview"
- Addleman, H., Dart, E., "BGP Architectures and Best Practices"
- Miller, K. "Security best practices in high-speed networks"

The workshop was attended by over 150 individuals over a 3-day period. The reviews were generally positive, but did note the need to improve:

- Lecture vs. hands on;
- Time allotment per talk, and overall (e.g. need for more breaks)
- Mechanisms to judge understanding of audience;

EPOC staff used the USC visualization labs as part of an expanded perfSONAR virtual training session in September for the LEARN community. LEARN was the recipient of a CC* award geared towards establishing a Science DMZ and DTN infrastructure between some of their smaller constituent institutions, and perfSONAR is an integral part of their plan to establish regular testing and performance analysis. The class covered the rationale behind using perfSONAR, examples of existing deployments, and basic command line configuration and troubleshooting. Attendees participated in hands-on exercises utilizing the virtual lab environment.

EPOC staff continue to work with the USC team to plan events that utilize a virtual lab environment. EPOC staff are discussing possible training sessions with the SoX regional network to take place later in 2021. The University of South Carolina team will also give a webinar to the CI engineering community in early 2021 to discuss the system, and ways it can enable education and research.

11.B Border Gateway Protocol (BGP) Training

EPOC has received a number of 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.

EPOC staff are working with partners from ESnet, GlobalNOC, and TACC to define an overall strategy for BGP best practices and training for the R&E community. Dart and Adleman presented "Network Strategies to Enable Data Intensive Science" [2] in May 2020 as part of the

CI Engineering Brown Bag webinar series. This talk focused on why proper BGP configurations are important for high throughput data transfers. Training on BGP in May [5] also highlighted the community need for this material.

Our work with BGP is complimented by the newly formed GNA-G / APAN Routing Working Group (<https://www.gna-g.net/join-working-group/gna-g-routing-wg/>), which Schopf and Addleman co-lead, along with Warrick Mitchell, AARNet. We envision using that platform to increase the message of the work that EPOC is doing.

In part to assess the current state of understanding, EPOC solicited feedback on network engineer BGP experience via a survey. The results showed:

- While many R&E network engineers are well versed in BGP configuration, they rarely review the route policy and routing tables on the routers in their network.
- Most respondents oversaw more than two wide area network connections to R&E networks and commodity providers, and many had four or more.
- Route policy was used extensively to control routing, and only 3% of respondents let BGP defaults dictate traffic flow
- Over 85% of the respondents said when they do review their routes, they do it manually, therefore highlighting the need for automated tools or scripts to help review route tables and BGP policy.

As a result, our Year 4 plans include focusing on these areas. We plan to create training materials and best practice documents, such as:

- The need for regular routing policy review and general maintenance of BGP, and where to find resources to help.
- The need for proper and maintained BGP configurations to keep science data flows on the more performant R&E networks instead of commodity/commercial networks.
- “About” documents for tools that can automate and help engineers regularly review routing tables, BGP configuration, and policy. These include
 - NetSage
 - RouteViews
 - Router Proxies/Looking Glasses
- Examples of how to adjust BGP to address poor routes, why they occur, and how to improve route effectiveness.

11.C PerfSONAR Training

In Year 3, we gave four presentations related to perfSONAR. These included training sessions for LEARN [13, 19], as part of the KIBER AHM talking about the Arcadia deployment [31] and as part of the larger May training workshop [4]. EPOC staff use perfSONAR extensively as part of the Roadside Assistance and Consulting project, as detailed in Section 6.

The Network Startup Research Center (NSRC), based at the University of Oregon, had previously worked with members of the perfSONAR community to create a series of instructional videos for perfSONAR in August 2016.

The existing NSRC perfSONAR video library was re-evaluated and content was marked for replacement or edits where needed. Remote work to record sessions on various topics in need of refreshing was started. These initial recordings for introductory perfSONAR topics were made by Chevalier and were sent to NSRC staff for editing and inclusion in the final project. With Chevalier's departure, this work is being handed back to the members of the perfSONAR Documentation and Training team. No further perfSONAR-focused work is expected in Year 4.

11.D Year 4 Training Plans

EPOC staff will approach Year 4 training via four focus areas:

- Working with the University of South Carolina Cybertraining Team to deliver virtual, and potentially physical, training efforts
- Continuing to work with NSRC and other community members on the Best Practice resources for network routing
- Participation with the perfSONAR training team to update content and deliver training when required
- Limited efforts to curate existing, and create new, written and video content that can be posted online to be viewed by the greater community when needed

The University of South Carolina Cybertraining Team continues to develop lessons that can be run virtually, or via live lectures, and targets emerging CI training opportunities for students and staff. EPOC is a collaborator with this group and serves to guide some of the materials of interest. We will continue to see opportunities to present alongside and help them get the word out via the large population (700+) of the CI engineering community through use of the CI-ENG mailing list.

The NSRC at the University of Oregon has created and curated numerous sets of materials to target emerging R&E networks in the correct policies and procedures to configure network routing. EPOC will continue to work with this group, and others, that are aligned on instructing proper methods to install, deploy, and operate networks, along with measuring the outcomes to ensure proper performance over time.

The perfSONAR group continues to release new software packages and has been updating training materials to keep pace with the development efforts. EPOC staff routinely communicate with perfSONAR and help to create materials and spread best common practices.

Lastly, EPOC will continue to explore ways to curate existing, and create new, videos and written content for the Fasterdata web site, as well as other venues. These efforts started during the pandemic and will continue during Year 4.

12. Data Privacy and Security

Basic security measures are being maintained and there were no security incidents to report for Year 3.

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 docs were updated for EPOC and are available online at <https://epoc.global/wp-content/uploads/2019/02/EPOC-Data-Privacy-Policy-21919.pdf>.

13. Reporting Against Deliverables

Table 5 lists the current deliverables and their status.

WBS #	Deliverables	Status
RA	ROADSIDE ASSISTANCE AND CONSULTATIONS	
RA.1	Adaptation of IN@IU, ESnet science engagement, and IRNC NOC PET process with expanded focus	Compl Y1
RA.2	Advertising roadside assistance and consulting	Ongoing
RA 3	Assist with ongoing RAs - Partners	
RA 3.1	iLight RA/C	
RA 3.1.1	C - IU-NOAA (24)	Started Y2Q1, Compl Y2Q2
RA 3.2	FRGP RA/C	
RA 3.2.1	C - Mines (4)	Started Y1, Compl Y2Q2
RA 3.2.2	C - Tribal (6)	Compl Y1
	See also RA 3.1.1 (24)	
RA 3.2.3	C - AIHEC (39)	Started Y2Q3, Compl Y2Q3
RA 3.2.4	C-ASU (49)	Started Y2Q3, Compl Y2Q3
RA 3.2.5	C - NOAA (63)	Started Y2Q3, Ongoing
RA 3.2.5	RA - NCAR (76)	Started Y2Q4, Ongoing
RA 3.2.6	C- FRGP (106)	Started Y3Q2, Compl Y3Q2
RA 3.2.7	C - FRGP (110)	Started Y3Q2, Compl Y3Q2
RA 3.2.8	C- CSU/NOAA (116)	Started Y3Q3, Compl Y3Q3
RA 3.3	LEARN/TACC RA/C	
RA 3.3.1	C - PVAMU (14)	Started Y1Q4, Compl Y2Q1
RA 3.3.2	C - TAMU (23)	Started Y2Q1, Compl Y2Q3
RA 3.3.3	C- Trinity (31)	Started Y2Q2, Compl Y2Q2
RA 3.3.4	C- PVAMU (36)	Started Y2Q2, Compl Y2Q3
RA 3.3.5	C- LEARN (62)	Started Y2Q3, Compl Y3Q4
RA 3.3.6	C - Baylor (66)	Started Y2Q4, Compl Y2Q4
RA 3.3.7	C- UTSA (69)	Started Y2Q4, Compl Y3Q2
RA 3.3.8	RA-Texas Tech (71)	Started Y3Q1, Compl Y3Q1
RA 3.3.9	RA-Arecibo (107) – formerly 4.57	Started Y2Q2, Ongoing
RA 3.3.10	C- TAMU (113)	Started Y2Q3, Ongoing
RA 3.3.11	C- TAMU (115)	Started Y2Q3, Compl Y3Q4
RA 3.3.12	C- LEARN (123)	Started Y2Q3, Compl Y3Q4
RA 3.3.13 (NEW)	C- TACC (135)	Started Y3Q4, Ongoing
RA 3.3.14 (NEW)	C- LEARN (145)	Started Y3Q4, Compl Y3Q4
RA 3.3.15 (NEW)	C- TAMUSA (155)	Started Y3Q4, Ongoing

RA 3.4	OARnet RA/C	
RA 3.4.1	C - UCinn (21)	Started Y2Q1, Compl Y2Q2
RA 3.4.2	C- OSC (32)	Started Y2Q2, Compl Y2Q3
RA 3.4.3	C- Kent (35)	Started Y2Q2, Compl Y3Q2
RA 3.4.4	C - Kent (51)	Started Y2Q3, Compl Y3Q3
RA 3.4.5	C - UHawaii-OSU (57)	Started Y2Q3, Compl Y3Q1
RA 3.4.6	C- UCinn (68)	Started Y2Q4, Compl Y2Q4
RA 3.4.7	C-Kent State (93)	Started Y3Q1, Compl Y3Q1
RA 3.4.7	C- Kent State (111)	Started Y3Q2, Compl Y3Q3
RA 3.4.8	C- U Miami (131)	Started Y3Q3, Compl Y3Q3
RA 3.4.9 (NEW)	C- KSU (144)	Started Y3Q4, Compl Y3Q4
RA 3.4.10 (NEW)	C- OARnet (151)	Started Y3Q4, Compl Y3Q4
RA 3.5	GPN RA/C	
RA 3.5.1	C - WSU (12)	Started Y1Q2, OBE Y1Q3
RA 3.5.2	C - UWisc-OneNet (25)	Started Y2Q1, Compl Y2Q1
RA 3.5.3	RA - Iowa-NCAR (27)	Started Y2Q1, Compl Y2Q3
RA 3.5.4	C- NDSU (48)	Started Y2Q3, OBE Y2Q4
RA 3.5.5	C- MSU Deep Dive (50)	Started Y2Q3, Compl Y3Q1
RA 3.5.6	RA SLU-Amazon (59)	Started Y2Q3, Compl Y3Q3
RA 3.5.7	C- U Missouri (61)	Started Y2Q3, Compl Y2Q3
RA 3.5.8	C- KanREN (65)	Started Y2Q3, Compl Y2Q4
RA 3.5.9	C- GPN (72)	Started Y2Q4, Compl Y3Q1
RA 3.4.10	C- OneNet (86)	Started Y3Q1, Compl Y3Q2
RA 3.4.11 (NEW)	C- UA (133)	Started Y3Q3, Compl Y3Q4
RA 3.4.12 (NEW)	C- OneNet (148)	Started Y3Q4, Compl Y3Q4
RA 3.4.13 (NEW)	C- SLU (153)	Started Y3Q4, Ongoing
RA 3.6	KINBER RA/C	
RA 3.6.1	C - F&M (17)	Started Y2Q1, Compl Y2Q1
RA 3.6.2	C - Duquesne (19)	Started Y2Q1, Compl Y2Q3
RA 3.6.3	C- Arcadia (29)	Started Y2Q2, Compl Y2Q2
RA 3.6.4	C- Penn State (42)	Started Y2Q2, Compl Y2Q2
RA 3.6.5	C- Duquesne (70)	Started Y2Q4, Compl Y2Q4
RA 3.6.6	C- Duquesne (73)	Started Y2Q4, Compl Y3Q3
RA 3.6.7	C- Penn State (87)	Started Y3Q1, Compl Y3Q2
RA 3.6.8	C- Lafayette (105)- Formerly 4.56	Started Y3Q2, Ongoing
RA 3.6.9	C- PSU (130)	Started Y3Q3, Compl Y3Q4
RA 3.6.10	C- Allegheny (132)	Started Y3Q3, Compl Y3Q3
RA 3.7	ESIP RA	Ongoing

RA 3.8	ICNWG RA	OBE
RA 3.9	IU GC RA	Ongoing
RA 3.10	UHawaii RA	Ongoing
RA 3.10.1	PANStarrs (1)	Compl Y1; 3x improvement
RA 3.10.2	U Hawaii (119)	Started Y3Q3, Ongoing
	See also RA 3.4.6	
RA 3.11	MWBDH RA	Ongoing
RA 3.12	OSN RA	OBE
RA 3.13	SoX RA	
RA 3.13.1	C- Vanderbilt (20) -prev. RA 4.11	Started Y2Q1, Compl Y2Q1
RA 3.13.2	C- U Southern Carolina (60) - prev. RA 4.32	Started Y2Q3, Compl Y2Q4
RA 3.13.3	C- ORNL (88)	Started Y3Q1, Compl Y3Q1
RA 3.14.4	C- SoX (129)	Starter Y3Q3, Compl Y3Q4
RA 3.15	Sun Corridor RA	
RA 3.15.1	Sun Corridor (81)	Started Y3Q1, Compl Y3Q3
RA 3.15.2	ASU (83)	Started Y3Q1, Compl Y3Q2
RA 3.15.3	ASU (99)	Started Y3Q2, Compl Y3Q2
RA 3.15.4	AZ Comm Colleges (118)	Started Y3Q3, Compl Y3Q4
RA 3.15.5 (NEW)	ASU (141)	Started Y3Q4, Compl Y3Q4
RA 3.15.6 (NEW)	ASU (143)	Started Y3Q4, Compl Y3Q4
RA 3.15.7 (NEW)	AMNH (160)	Started Y3Q4, Ongoing
RA 3.16	CENIC RA	
RA 3.16.1	Allen Inst (74) Formerly RA 4.35	Started Y2Q4, Compl Y3Q4
RA 3.16.2	LBNL (101) Formerly RA 4.52	Started Y3Q2, Compl Y3Q4
RA 3.16.3	Allen Inst (104) Formerly RA 4.55	Started Y3Q2, Compl Y3Q4
RA 3.16.4	Allen Inst (117)	Started Y3Q3, Ongoing
RA 3.16.5 (NEW)	UAF (142)	Started Y3Q4, Ongoing
RA 3.16.6 (NEW)	USC (146)	Started Y3Q6, Compl Y3Q4
RA 4	Other RA/C	
RA 4.1	LHC Pakistan (2)	Compl Y1; 10x improvement
RA 4.2	C - New York University School of Medicine (5)	Compl Y1
RA 4.3	C – AMNH (7)	Started Y1, Compl Y2Q2
RA 4.4	C- UF (8)	Compl Y1
RA 4.5	C- LSU Health (9)	Started Y2Q1, Compl Y2Q1
RA 4.6	C- SANReN (10)	Started Y2Q1, OBE Y2Q3
RA 4.7	C- PNNL (11)	Started Y2Q1, Y2Q4
RA 4.8	C - Compute Canada (13)	Compl Y1
RA 4.9	C- UC Merced (15)	Started Y2Q1, Compl Y2Q3

RA 4.10	C - LSU Health Deep Dive Templates (18)	Started Y2Q1, Compl Y2Q1
RA 4.11	Now RA 3.12.1	
RA 4.12	C - UWisc - MichSt (26)	Started Y2Q1, OBE Y2Q3
RA 4.13	C - UC Merced (28)	Started Y2Q1, Compl Y2Q1
RA 4.14	C- SANReN(30)	Started Y2Q2, Compl Y2Q2
RA 4.15	C- AMNH (33)	Started Y2Q2, Compl Y2Q3
RA 4.16	C- U Mich (34)	Started Y2Q2, Compl Y2Q4
RA 4.17	C- UNCG (37)	Started Y2Q2, Compl Y2Q4
RA 4.18	C- U Mich (38)	Started Y2Q3, Compl Y2Q3
RA 4.19	C- AAMU (40)	Started Y2Q3, OBE Y2Q3
RA 4.20	C- UC Davis (41)	Started Y2Q3, OBE Y2Q3
RA 4.21	C-MGHPC (43)	Started Y2Q3, Compl Y2Q3
RA 4.22	C-AMNH (44)	Started Y2Q3, OBE Y2Q4
RA 4.23	C - Wayne (45)	Started Y2Q3, Compl Y2Q3
RA 4.24	C- U Wisc (46)	Started Y2Q3, Compl Y2Q3
RA 4.25	C-UCentral FL (47)	Started Y2Q3, Compl Y2Q3
RA 4.26	C- U Montana (52)	Started Y2Q3, Compl Y2Q4
RA 4.27	C- CalTech (53)	Started Y2Q3, Compl Y2Q4
RA 4.28	C-Globus (54)	Started Y2Q3, Compl Y2Q3
RA 4.29	C- U Montana (55)	Started Y2Q3, Compl Y2Q4
RA 4.30	C- U Montana (56)	Started Y2Q3, Compl Y2Q4
RA 4.31	C-OSHEAN (58)	Started Y2Q3, OBE Y2Q4
RA 4.32	Now RA 3.12.2	
RA 4.33	C-U Southern California DMZ (64)	Started Y2Q3, Compl Y2Q4
RA 4.34	C - VA (67)	Started Y2Q4, Compl Y3Q1
RA 4.35	Now RA 3.16	
RA 4.36	C - Reed (75)	Started Y2Q4, Compl Y3Q4
RA 4.37	C - Compute Canada (77)	Started Y2Q4, Compl Y3Q1
RA 4.38	C - MIT (79)	Started Y2Q4, Compl Y2Q4
RA 4.39	C - UCentral FL (80)	Started Y2Q4, Compl Y3Q1
RA 4.40	SanREN (84)	Started Y3Q1, Compl Y3Q1
RA 4.41	MDREN (85)	Started Y3Q1, Compl Y3Q1
RA 4.42	OSHEAN (89)	Started Y3Q1, Compl Y3Q1
RA 4.43	UCSC-ASTAR (90)	Started Y3Q1, Compl Y3Q2
RA 4.44	RNP (91)	Started Y3Q1, Compl Y3Q2
RA 4.45	UCF (92)	Started Y3Q1, Compl Y3Q3
RA 4.46	NLM/ASGC (94)	Started Y3Q1, Compl Y3Q3
RA 4.47	CCNY/JGN (95)	Started Y3Q1, Compl Y3Q4

RA 4.48	Yale (96)	Started Y3Q1, Compl Y3Q2
RA 4.49	UCF (97)	Started Y3Q1, Compl Y3Q2
RA 4.50	126.com (98)	Started Y3Q2, Compl Y3Q2
RA 4.51	EBI (100)	Started Y3Q2, Compl Y3Q3
RA 4.52	Now RA 3.16.3	
RA 4.53	USDA (102)	Started Y3Q2, Compl Y3Q4
RA 4.54	VRO/LSST (103)	Started Y3Q2, Compl Y3Q2
RA 4.55	Now RA 3.16.3	
RA 4.56	Now RA 3.6.8	
RA 4.57	Now RA 3.3.10	
RA 4.58	KAUST (108)	Started Y3Q2, Compl Y3Q4
RA 4.59	NIH (109)	Started Y3Q2, Compl Y3Q3
RA 4.60	U South Carolina (112)	Started Y3Q2, Compl Y3Q4
RA 4.61	NYU (114)	Started Y3Q3, Ongoing
RA 4.62	Yale (120)	Started Y3Q3, Compl Y3Q4
RA 4.63	EHT (121)	Started Y3Q3, Ongoing
RA 4.64	SUNY Medical (122)	Started Y3Q3, Compl Y3Q4
RA 4.65	Oregon State (124)	Started Y3Q3, Compl Y3Q4
RA 4.66	WSU (125)	Started Y3Q3, Compl Y3Q4
RA 4.67	KAUST (126)	Started Y3Q3, Compl Y3Q3
RA 4.68	HHU (127)	Started Y3Q3, Compl Y3Q3
RA 4.69	EBI/ NYGC (128)	Started Y3Q3, Compl Y3Q4
RA 4.70	U Miss (134)	Started Y3Q3, Compl Y3Q4
RA 4.71 (NEW)	Yale (136)	Started Y3Q4, Compl Y3Q4
RA 4.72 (NEW)	SingaREN (137)	Started Y3Q4, Compl Y3Q4
RA 4.73 (NEW)	SingaREN (138)	Started Y3Q4, Compl Y3Q4
RA 4.74 (NEW)	PPPL (139)	Started Y3Q4, Compl Y3Q4
RA 4.75 (NEW)	SingaREN (140)	Started Y3Q4, Compl Y3Q4
RA 4.76 (NEW)	SARAO (147)	Started Y3Q4, Ongoing
RA 4.77 (NEW)	Arecibo (149)	Started Y3Q4, Ongoing
RA 4.78 (NEW)	UNC-C(150)	Started Y3Q4, Compl Y3Q4
RA 4.79 (NEW)	BU (152)	Started Y3Q4, Ongoing
RA 4.80 (NEW)	NyserNet (154)	Started Y3Q4, Ongoing
RA 4.81 (NEW)	HARNET (156)	Started Y3Q4, Compl Y3Q4
RA 4.82 (NEW)	Syracuse (157)	Started Y3Q4, Ongoing
RA 4.83 (NEW)	NYULH (158)	Started Y3Q4, Ongoing
RA 4.84 (NEW)	Yale (159)	Started Y3Q4, Ongoing
RA 4.85 (NEW)	AMNH (160)	Started Y3Q4, Ongoing

RA 4.86 (NEW)	RPI (161)	Started Y3Q4, Ongoing
DD	DEEP DIVE	
DD.1	Adaptation of ESnet facility deep dive process for use with applications	Compl Y1
DD.2	Over project period, goal is to offer at least 2 deep dives per regional partner	Ongoing
DD.2.1	iLight Deep Dives	Ongoing
DD 2.1.1	Purdue University	Compl - Event Y2Q1, report Y2Q3
DD.2.2	FRGP Deep Dives	Ongoing
DD 2.2.1	NOAA and NASA Deep Dive (with Training)	Compl Y1
DD 2.2.2	Now DD 2.13.1	
DD 2.2.3	N Arizona	OBE (COVID)
DD 2.2.4 (NEW)	NOAA Virtual	Virtual started Y3Q4
DD 2.3	LEARN Deep Dives	Ongoing
DD 2.3.1	Trinity University	Compl - Event Y2Q1, report Y2Q3
DD 2.3.2	Baylor	Compl - Event Y2Q4, report Y3Q4
DD 2.3.2	TAMUSA	Virtual started Y3Q4
DD 2.4	OARnet Deep Dives	Ongoing
DD 2.4.1	University of Cincinnati	Compl - Event Y2Q1, report Y2Q3
DD 2.5	GPN Deep Dives	Ongoing
DD 2.5.1	Training - KSU Agronomy	Compl - Event Y2Q1, report Y2Q3
DD 2.5.2	University South Dakota	Virtual started Y3Q4
DD 2.6	KINBER Deep Dives	Ongoing
DD.2.6.1	Arcadia Bioinformatics (with training)	Compl - Event Y2Q1, report Y2Q2
DD 2.7	ESIP DD	Ongoing
DD 2.8	ICNWG DD	OBE
DD 2.9	IU GC RDD	Ongoing
DD 2.10	UHawaii DD	Ongoing
DD 2.11	MWBDH DD	Ongoing
DD 2.12	OSN DD	OBE
DD 2.13 (NEW)	SCN Deep Dive	
DD 2.13.1 (NEW)	Arizona State/Sun Corridor	On hold (COVID)
DD.3	Other Deep Dives	Ongoing
DD.3.1	QUILT/University Maryland (with Training)	Compl Y1
DD.3.2	University of Wisconsin	Event Y2Q1, report expected Y3Q1
DD 3.3	PEARC'19	Compl (no report)
DD 3.4	Oregon State Univ	On hold (COVID)
DD 3.5	Quilt Briefing	Compl Y2Q4
DD 3.6	University Central Florida	Virtual started Y3Q4

NS	NETSAGE	NOTE: Renumbering took place in Y2Q2
NS.1	NetSage prototypes for regional partners	Ongoing
NS1.1	NetSage for iLight	Ongoing
NS 1.1.1	SNMP for iLight	Not needed
NS 1.1.2	Flow for iLight	Initial Y2Q1, Ongoing support
NS 1.2	NetSage for FRGP	Ongoing
NS 1.2.1	SNMP for FRGP	Discussion Ongoing
NS 1.2.2	Flow for FRGP	Initial Y2Q4, Ongoing support
NS 1.2.3	Tstat for NOAA	Compl -Deployed Y2Q1, OBE
NS 1.2.4	Tstat for NCAR	Compl Y2Q2
NS 1.3	NetSage for LEARN/TACC	Ongoing
NS 1.3.1	SNMP for LEARN	Discussion Year 3
NS 1.3.2	Flow for LEARN	Discussion Year 3
NS 1.3.3	Tstat on TACC archives	Compl Y1, updated Y2
NS 1.3.4	SNMP for TACC	Discussion ongoing
NS 1.3.5	Flow for TACC	Initial Y3Q2, Ongoing support
NS 1.4	NetSage for OARnet	Ongoing
NS 1.4.1	SNMP for OARnet	OBE
NS 1.4.2	Flow for OARnet	OBE
NS 1.5	NetSage for GPN	Ongoing
NS 1.5.1	SNMP for GPN	Initial Y1, Ongoing support
NS 1.5.2	Flow for GPN	Initial Y2Q2, Ongoing support
NS 1.6	NetSage for KINBER	Ongoing
NS 1.6.1	SNMP for KINBER	OBE
NS 1.6.2	Flow for KINBER	Initial Y2Q3, Ongoing support
NS 1.7	NetSage for Sun Corridor Network (SCN)	Ongoing
NS 1.7.1	SNMP for SCN	Discussion Ongoing
NS 1.7.2	Flow for SCN	Initial Y3Q4, ongoing support
NS 1.8	NetSage for SoX	Ongoing
NS 1.8.1	SNMP for SoX	Discussion Ongoing
NS 1.8.2	Flow for SoX	Initial Y3Q3, Ongoing support
NS 1.9	NetSage for CENIC/ Pacific Wave	Ongoing
NS 1.9.1	SNMP for CENIC	Transition from IRNC Y3Q4
NS 1.9.2	Flow for CENIC	Transition from IRNC Y3Q4

NS 2	NetSage deployments related to other partners	Ongoing
NS 2.1	University Hawaii	Ongoing
NS 2.1.1	Tstat on Astronomy Archive	Compl Y2Q1, ongoing support
NS 3	Adaptation of NetSage analysis for network disturbance detection	Ongoing
NS 3.1	Examine COVID-related data transfer performance	Started Y3Q1, Ongoing
D/P	Data Mobility Exhibition	Note: Restructured Y3Q1
DP 1	Data Mobility Exhibition Setup	Ongoing
DP 2	Data Mobility Exhibition Support	Ongoing
MS	MANAGED SERVICE	Note: Numbering reworked Y2Q3
MS 1	Define Managed Services	Ongoing
MS 1.1	Define perfSONAR Managed Service (PS MS)	Started Y1, Ongoing
MS 1.2	Define DMZ Managed Service (DMZ MS)	Delayed (COVID)
MS 1.3	Define Data Transfer Managed Service (DT MS)	Delayed (COVID)
MS 1.4	Tester Managed Service	Definition Compl Y2
MS 1.5	Portal Prototype Development	Ongoing
MS2	MS deployments	Ongoing
MS 2.1	iLight MS	TBD
MS 2.2	FRGP MS	TBD
MS 2.2.1	PS MS for Tribal Colleges	Completed Y2
MS 2.3	LEARN MS	Underway Y2
MS 2.3.1	LEARN DMZ MS	Delayed (COVID)
MS 2.3.2	LEARN DT MS	Delayed (COVID)
MS 2.3.3	Portal with Baylor (LEARN)	On hold
MS 2.4	OARnet MS	TBD
MS 2.4.1	OARnet DT MS	OBE
MS 2.4.2	Testset Loan to Kent State	OBE
MS 2.5	GPN MS	TBD
MS 2.5.1	GPN and KanREN DT MS	Delayed (COVID)
MS 2.5.2	Portal wirth GPN member	On hold
MS 2.6	KINBER MS	Started Y1, Ongoing
MS 2.6.1	KINBER and Arcadia PS MS	Completed Y3Q4
MS 2.6.2	Testset Loan to Duquesne	OBE
MS 2.7	Other MS Deployments	Ongoing

MS 2.7.1	U Hawaii Astronomy Data Portal MS	On hold
T	TRAINING	
T 1	Set up public site for training materials	Compl Y1
T 2	Technical training	Ongoing
T 2.1	SOX - perfSONAR	Compl Y1
T 2.2	GPN LCI - perfSONAR, DMZ	Compl Y2Q1
T2.3	LEARN - PS, DMZ, DTN, Security	Compl Y2Q1
T 2.4	NWT Star/FRGP - PS, DMZ, DTN, Security	Compl Y2Q2
T 2.5	CyberTraining w/USC -PS, DMZ, DTN, Engagement	Compl Y2Q2
T 2.6	Managed Service PS with KINBER, Arcadia	Delayed (COVID)
T 2.7	CyberTraining w/USC - BGP, PS, DMZ	Compl Y3Q1 (virtual)
T 2.8	CyberTraining w/USC - BGP, PS, DMZ	Delayed Y3Q3 (virtual)
T 2.9	PS Training for LEARN	Compl Y3Q2 (virtual)
T 2.10	CyberTraining w/USC - BGP, PS, DMZ for SoX	Compl Y3Q4 (virtual)
T 2.10	LEARN PS train with lab	Compl Y3Q4 (virtual)
T 3	Deep Dive training	Ongoing
T3.1	NOAA DD Training	Compl Y1
T 3.2	QUILT DD Training	Compl Y1
T 3.3	KINBER DD Training	Compl Y2Q1
T 3.4	GPN DD Training	Compl Y2Q1
T 3.5	PEARC DD Training	Compl Y2Q2
T 3.6	Quilt DD Training	OBE (overview on request)
T 3.7	DD Training 6	TBD (COVID)
T 3.8	DD Training 7	TBD (COVID)
T 4	Other Related General Activities	TBD as requested by community
T 4.1	Finding Researchers	iLight - Compl Y2Q1
T 4.2	Data Mobility Expo	Compl Y2Q2
T 4.3	BGP BOF at I2 TechEx	Compl Y2Q3
T 4.4	BGP BOF at I2 Global Summit	Delayed (COVID)
T 4.5	PS NSRC Updates	Compl Y3Q4
T 4.6	10G Easy DTN Video	Started Y3Q1, ongoing
T 5	Reworking Training during COVID	Ongoing
T 5.1	New plan development	Ongoing