

Engagement and Performance Operations Center (EPOC)

(Formerly known as ReSEC)

NSF Award #1826994

Year 3 Quarter 2

1 July 2020 through 30 September 2020

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 Quarter 2 highlights include increasing the use of the Data Mobility Exhibition/Portal to enable an institution baseline for moving a TeraByte of data in an hour, additional deployments of NetSage with the TACC network and Southern Crossroads (SoX), and closing fourteen Consultations.

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 Mover Exhibition and associated work with our simplified portal to check transfer times against known “good” end points;
- Provision of Managed Services via support through the IU GlobalNOC 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 in Year 3 Quarter 2. 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 Quarter 2, 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, Data Mobility Exhibition, Roadside Assistance
- 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 0.5FTE of developer support for the NetSage deployments. No staff changes were made during this quarter.

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.

Our next formal Team All Hands meeting will take place in early 2021 and will likely be virtual.

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. Remote participation in Quarter 2 meetings included:

- Schopf attended the PEARC'20 meeting, July 27-31, <https://pearc.acm.org/pearc20/>, where she learned about other US domestic CI developments that may play a role in our continuing work.
- Moynihan attended the TICAL 2020 conference, virtual, Aug 31-Sept 4, 2020, <https://tical2020.redclara.net/>. He attended and participated in sessions on supporting the Vera Rubin Observatory, the Bella project, and on supporting COVID-19 research collaborations.
- Schopf and Addleman attended the GNA-G Virtual Meeting on September 14-15, 2020. They discussed participating with several new working groups will likely form as they become more established.
- Moynihan attended the LHCONE Virtual meeting on September 16-17, 2020 <https://indico.cern.ch/event/932306/>. He attended and participated in updates on the LHCONE network and on future infrastructure needs.
- Southworth and Chevalier presented for the LEARN perfSonar Virtual Training on September 21-22, 2020. They presented an introduction to perfSONAR course along with hands-on virtual lab activities provided by the University of South Carolina, as well as an overview of Science DMZ.
- Addleman attended the NSF Virtual CyberSecurity Summit on September 22-24, 2020 <https://www.trustedci.org/2020-nsf-summit>. During the conference he attended sessions that included updates to cybersecurity regulations, disinformation tracking on social media, socializing cybersecurity in an organization, and a talk on the trustworthiness of data.
- Schopf, Addleman, Miller, and Zurawski attended the Quilt Fall Virtual Meeting on September 30 - October 1, 2020 <https://www.thequilt.net/public-event/2020-quilt-virtual-fall-member-meeting/>. They led a 45 minute overview that discussed EPOC at a high level, with more detail about the Deep Dives and DME/Portal activities.

4. Presentations and Publications

For Year 3 Quarter 2, the EPOC presentations, invited posters, and publications are listed here and referred to throughout the report with the reference number listed:

1. Southworth, Douglas, "Whose Line is it Anyway?: Problem Solving in Complex Networks", Trusted CI Monthly Webinar Series, July 20, video online at <https://www.youtube.com/watch?v=TKH-6CNO2mc>
2. Southworth, Douglas, "Introduction to perfSONAR", LEARN Technical Advisory Group webinar, August 10.

3. Zurawski, Jason, "Identifying Campus Infrastructure Needs", Quilt CC* Seminar, August 13, video online at <https://youtu.be/EqNvKQWwYuQ>
4. Miller, K., "Data Movement Exhibition" Invited Talk, CI Engineering Lunch and Learn on August 14, 2020, Talk available online at: <https://youtu.be/CmHGu9cG0ww>
5. Robb III, George B., Pejovic, Goran, Shiflett, Charles, "Data Transfer Node Design Patterns" Invited Talk, CI Engineering Lunch and Learn on August 28, 2020, Slides available online at:
https://drive.google.com/file/d/1ABSe4OalAROTBi0ltM5VXcLORQ2yaNa_/view?usp=sharing
6. Schopf, J, "NSF EPOC and University of Cincinnati Campus-Wide Deep Dive", Invited Talk jointly with J. Combs, University Cincinnati, CaRCC Data Facing Track Call, September 1, 2020, video at https://youtu.be/jhg5R9_VdQY
7. Miller, K. "Data Movement Exhibition", Invited Talk, Quilt Managed Services Working Group on September 8th, 2020, video online at <https://youtu.be/wuZrP2Rg1IA>
8. Southworth, D. and Chevalier, S. "perfSonar Workshop", Training, LEARN Virtual Training on September 21-22, 2020, Slides available online at:
https://drive.google.com/drive/folders/1jMAG6uw_-ESvATcb895xg-JWc6b7oxjk
9. Schopf, J. "Moving Data Faster with the Engagement and Performance Operations Center (EPOC)", Invited Talk, Quilt Fall Virtual Meeting on September 30, 2020, Slides available online at:
<https://drive.google.com/file/d/1OsktyJniXVgRkPJP8apkP3jXFKMpl0M3/view?usp=sharing>
10. Miller, K. "Data Movement Exhibition", Invited Talk, Quilt Fall Virtual Meeting on September 30, 2020, Slides available online at:
<https://drive.google.com/file/d/1NywUeYH4xvImiwb4pqBvpxl9sPafTkl/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.

Our first three additional Regional Networking Partners are Southern Crossroads (SoX), the regional network for the state of Arizona, Sun Corridor, the regional network for much of the southeastern part of the US including parts of Alabama, Georgia, South Carolina, and Tennessee, and 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. We are currently investigating if more formal documentation would be helpful.

5.B Regional Networking Partners

EPOC is partnered with the six regional network operators.

- **Front Range GigaPop (FRGP)** is the regional collaboration of networks that cover the western states, including Colorado, Wyoming, Arizona, Idaho, Utah, and New Mexico.
 - Roadside Assistance: #76
 - Consultations: 63, 106, 110
 - 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.
- **iLight** is the regional network for Indiana.
 - NetSage: <http://ilight.netsage.global>, deployment May 2019, working on ASN splitting with SWIP.
- **The Great Plains Network (GPN)** is the regional network that serves North Dakota, South Dakota, Nebraska, Iowa, Minnesota, Kansas, Missouri, and Arkansas.
 - Roadside Assistance: #59
 - Consultations: # 86
 - Deep Dives: In planning with University of South Dakota, 2021 timeframe.
 - NetSage: <http://gpn.netsage.global>, SNMP deployment completed 2019, Flow data deployment using containers deployed this quarter.
- **The Keystone Initiative for Network Based Education and Research (KINBER)** is the regional network for Pennsylvania.
 - Consultations: #73, 87, 105
 - NetSage: <http://kinber.netsage.global> Deployment November 2020, working on ASN extensions with SWIP.

- Managed Services: We continue to work with Arcadia University and KINBER on deploying a perfSONAR managed service at Arcadia.
- **The Lonestar Education and Research Network (LEARN)** is the regional network for Texas.
 - Meetings: Presentation on perfSONAR at the LEARN Technical Advisory Group [2], perfSONAR virtual training for LEARN members in September [8]
 - Roadside Assistance:
 - Consultations: #62, 69, 113 (Note: we include these for both LEARN and TACC)
 - Deep Dives: January 2020 with LEARN partner Baylor, report expected in Year 3
 - NetSage: Continued discussion for deployment.
 - DME/Portal: LEARN-member Baylor University investigating use of Portal service
- **The Ohio Academic Resources Network (OARnet)** is the regional network for Ohio.
 - Consultations: #35, 51, 111
 - NetSage: Continued discussion for deployment.
- **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: Invitation to present at the SoX All Hands Meeting in October.
 - NetSage: Deployment of NetSage for Flow data went live in July 2020, in discussions for SNMP deployment.
- **Sun Corridor Network (SCN)** is the regional network for the state of Arizona
 - Consultations: 81, 83, 99, 113
 - NetSage: Deployment in discussion.
 - Deep Dive: Discussions for a Deep Dive to take place in 2021 are ongoing.
- **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
 - Consultations: #62, 69, 113 (Note: we include these for both LEARN and TACC)
 - NetSage: Deployment for TACC completed in July 2020. Archive deployment in production since 2019.

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 tracking many of their mailing lists and participating in several of them actively. She gave a talk as part of the Data Facing Track Call [6], and offered to give a talk on the CC* solicitation when it was live for the Emerging centers call.

- **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 [1]. 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.
- **Internet2** supports solving common technology challenges for their over 200 educational, research, and community members. Schopf is the main contact for this organization. Internet2 is holding their TechEx replacement, Tech Expo, and we are trying to find out how to participate. In general, we are currently seeking additional ways to work with Internet2 in the absence of face-to-face meetings.
- **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 staff are involved in the ongoing Quilt-run seminars related to the CC* program [3, 7], and participated in the Virtual Fall Quilt Member Meeting [9, 10].
- **The Science Gateway Community Institute (SGCI)** provides best practice recommendations and support for scientists building and using data portals. In Quarter 2, SGCI updated their partnership program based on community input, and we affirmed our status as project partners for 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.

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. The July meeting was virtual. 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.
- **The Midwest Big Data Hub (MBDH)** supports the use of data for a variety of applications and end users across twelve states. Southworth planned to attend the All Hubs Meeting, a conference which includes all of the regional hubs, in June 2020. This meeting was cancelled due to COVID related travel restrictions. No updated meeting date has been set at this time.

5.E 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. This quarter, EPOC staff gave two talks in support of this series [4,5]. To date, 100 webinars have been uploaded to the EPOC YouTube channel as of September 30, 2020, and available online at: <https://www.youtube.com/channel/UChIaulc1bccif1Dz4cfZl0w>.

6. Roadside Assistance and Consulting

A key aspect of the EPOC project is the operations center and process pipeline for immediate help, referred to as Roadside Assistance and Consulting, which helps collaborators when data sharing failures occur. EPOC coordinates with the multiple domains and organizations involved to achieve a resolution. More information about the Roadside Assistance and Consulting process is available at: <https://epoc.global/wp-content/uploads/2019/04/Roadside-Assistance-Description.pdf>. Hans Addleman is the lead for this effort.

6.B Roadside Assistance Cases

In Year 3 Quarter 2, we had four ongoing Roadside Assistance cases:

- **59 - Saint Louis University transfers to Amazon S3 Cloud:** A biologist at Saint Louis University (SLU) experiencing very poor file transfer performance 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 from there to save time. EPOC reached out to SLU IT staff, but the university has outsourced much of its network support to an external contractor, which made finding the correct engineer to contact more challenging. 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 continues to engage with the SLU IT group, Missouri Research and Education Network (MOREnet), and the Great Plains Network (GPN) to debug this issue. During Quarter 1, EPOC staff met with SLU's Director of Infrastructure who confirmed that the network was in the process of being upgraded and asked us to re-engage with SLU in early 2021.

- **76 - NCAR Multicast Performance:** NCAR/UCAR 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 later in 2021.
- **94 - National Library of Medicine to Academia Sinica Grid Computing Center (ASGC):** 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. These routing discussions are ongoing and EPOC staff are continuing to lend support where necessary.
- **100 - European Bioinformatics Institute (EBI) to Partners Healthcare Systems (PHS):** 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 were able to determine that the bottleneck was not at EBI or on the trans-Atlantic circuits. We have now started a conversation with engineers from the Northern Crossing (NOX) exchange to learn more about the set-up of the PHS network and to see if they can introduce us to engineers at PHS who could help us determine the issue.

6.C Consultation Cases

In Year 3 Quarter 2, we had fourteen completed and eighteen ongoing consultations. The primary topics were assistance with 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.
- **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.
- **83 - Arizona State University Sustainable Cities Network (ASU SCN)**: EPOC reviewed the SCN CC* proposal with ASU staff and later discussed ScienceDMZ architecture for small campuses.
- **86 - OneNet**: EPOC staff spent time with OneNet staff reviewing questions they received from an NSF review panel about switch buffering, network hardware, and their proposed ScienceDMZ. OneNet received NSF funding of the proposal: <https://onenet.net/si-more-offn-award-connects-additional-institutions-to-oklahomas-science-dmz/>.
- **87 - Penn State to the Advanced Light Source at Berkeley (ALS)**: Penn State reached out to an EPOC team member directly as part of their set up to support computed tomography (CT) imaging at the ALS since Argonne will be closed for an upgrade next year. Penn State staff requested estimates for expected file transfer rates between Penn State and ALS Berkeley. EPOC staff provided baseline iperf tests between the borders.
- **91 - Rede Nacional de Ensino e Pesquisa (RNP)**: EPOC engineers consulted with the Brazilian National Research and Education Network (NREN), RNP, regarding specifications for 100G capable data transfer nodes and shared known performance results for various configurations.
- **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.
- **103 - Rubin Observatory:** EPOC exchanged emails with a Rubin engineer about using the perfSONAR CLI and API. EPOC also directed the engineer to the perfSONAR users group mailing list for more advanced discussion.
- **106 - Front Range GigaPop (FRGP):** Engineers at FRGP inquired about the role of buffer sizes and network protocols in supporting science, as well as recent developments that may reduce buffer importance and overall equipment cost for smaller schools. EPOC shared the best current operational practices around the new protocols and buffering.
- **110 - Front Range GigaPop (FRGP):** FRGP engineers asked for a link to a transmission control protocol (TCP) BBR slide deck referenced during a Cyber Infrastructure Data Transfer Node (DTN) video. EPOC provided links to various talks on the BBR protocol.

Ongoing Consultations in Quarter 2 included:

- **51 - Kent State University (KSU):** KSU engineers inquired about borrowing the EPOC Viavi network tester in early 2020. They will follow up when their deployment is ready for testing.
- **62 - Lonestar Education and Research Network (LEARN):** LEARN engineers are asking for help designing a low cost but powerful DTN. They are currently reviewing the information EPOC staff has provided. They have asked to engage again in Quarter 3.
- **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 will be presenting virtually on the process during a NOAA meeting in Quarter 3.
- **73 - Duquesne University:** Engineers from Duquesne and EPOC staff discussed testing methodologies for equipment loaned by vendors for evaluation of the Duquesne ScienceDMZ use cases. During this quarter EPOC sent the Viavi network tester to Duquesne. Engineers will be using it to stress test vendor hardware in their lab before choosing a vendor for their 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. They have asked to re-engage in 2021.
- **75 - Reed College (RC):** A researcher at RC 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 was submitted in August, and our conversation is ongoing.
- **81 - Sun Corridor:** Staff from Sun Corridor, the regional R&E network for Arizona, and EPOC met via Zoom to review their ScienceDMZ CC* award. Topics discussed included architecture, DTN, perfSONAR, intrusion detection systems including Zeek, and firewalls. EPOC will follow up on the status of their proposal in Quarter 3.
- **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. A follow-up call is pending.
- **95 - City College of New York (CCNY) to Japanese Gigabit Network (JGN):** 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 have engaged engineers at CCNY to set up perfSONAR tests between the sites. At this time, testing has not replicated the packet loss. Several configuration changes are taking place at CCNY and they have engaged their campus engineers to look at possible firewall issues. Troubleshooting is ongoing.
- **101 - Lawrence Berkeley National Laboratory (LBNL):** EPOC staff and perfSONAR developers are helping a user at LBNL troubleshoot running perfSONAR inside of a docker container on a laptop. This allows perfSONAR testing to happen close to the end point that was experiencing performance issues.
- **104 - Allen Institute-Princeton:** An engineer at the Allen Institute reported performance issues with a science data transfer to a commercial cloud provider via Princeton University. EPOC engineers are working with both Princeton and Allen Institute to diagnose the issue. Princeton reports they are in the middle of network upgrades and will report back in Quarter 3.
- **105 - Lafayette College:** The manager of research and high performance computing at Lafayette College reached out for information about Science DMZs, DTN design, and data architecture best practices. This consultation is ongoing.
- **107 - Arecibo Observatory:** UCF asked that we consult with the staff at Arecibo Observatory about moving multiple petabytes of data off site to a Microsoft Azure Cloud storage instance and a subset of data to the University of Central Florida's Advanced Research Computing Center (ARCC) cluster for data processing. Discussion via email is ongoing and we will have a zoom meeting in October. Initial results show that they are getting on average 224mbps from the Arecibo facility to UCF.

- **108 - King Abdullah University of Science and Technology (KAUST):** A KAUST engineer reached out for information on 100G capable data transfer nodes and how best to support their researchers. This consultation is ongoing.
- **109 - National Human Genome Research Institute (NHGRI) NIH:** Globus connected EPOC staff to a program analyst at NHGRI in the process of setting up a Globus installation and Science DMZ. EPOC shared best practices, slide decks, and links during a video conference and follow-up email. This consultation is ongoing.
- **111 - Kent State:** A doctoral student at Kent state has questions about using perfSONAR data for research. EPOC is working with the researcher to understand available data sources and access processes.
- **112 - University of South Carolina (USC):** Staff at USC got a CC* grant and are in the process of deploying data transfer and perfSONAR nodes. Engineers from USC connected with EPOC to learn about routing and Science DMZ best practices. A zoom meeting was had and the consultation is ongoing via email.
- **113 - Texas A&M University (TAMU):** NetSage partners found a 7 day transfer between Texas Advanced Computing Center (TACC) and TAMU as part of their release testing that transferred over 24 TB of data at a rate under 500 Mbps. EPOC engineers are working with TACC and TAMU staff to determine the end points and next steps. The causes being considered are poor routing, poor end host, or use of the wrong protocol for the type of data being transferred (or all 3).

When a Roadside Assistance Case or Consultation does not respond to multiple attempts to move it forward over a 6+ week time frame, we consider those Cases Overtaken by Events (OBE). In Year 3 Quarter 2 we had 1 OBE Roadside Assistance Case, which started on June 30th, 2020 and was considered OBE on September 25th, 2020, due to communication from the consultee indicating that they were putting the project involved on indefinite hold.

- **90 - University of California Santa Cruz to A*STAR:** 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.

6.D Metrics

Table 1: A summary of Year 3 Quarter 2 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	Kent State Univ	OARnet	Cons	9/24/19	8/20/20	DTN	Eng	Infra	R	S
51	Kent State University	OARnet	Cons	11/1/19		Tester, PS	Eng	Infra	R	S
59	St Louis University	GPN	RA	12/11/19		Trans Perf	Scientist	Bio	E	S
62	LEARN	LEARN	Cons	12/13/19		DTN	Eng	Infra	O	-
63	NOAA	FRGP	Cons	12/16/19		DD	Eng	Infra	O	-
69	UT San Antonio	LEARN	Cons	1/21/20	8/20/20	DMZ, Security	Eng	Infra	R	L
73	Duquesne University	KINBER	Cons	2/2/20		Tester	Eng	Infra	E	S
74	Allen Institute	(CI LIST)	Cons	2/28/20		DMZ, DTN	Eng	Bio	R	S/M
75	Reed College		Cons	3/5/20		DMZ, DTN, Globus	Eng	Bio	E	S
76	NCAR	FRGP	RA	3/6/20		Trans Perf	Eng	Climate	R	L
81	Sun Corridor	Sun Corridor	Cons	4/3/20		DMZ	Eng	Infra	E	L
83	Arizona State University	SCN	Cons	4/6/20	8/20/20	DMZ, Security	Eng	Infra	E	S
86	OneNet	GPN	Cons	4/29/20	9/3/20	DMZ	Eng	Infra	E	
87	Penn State	KINBER	Cons	4/30/20	8/11/20	Trans Perf	Scientist	Bio	E	L

90	UC Santa Cruz	TP	RA	5/20/20	9/16/20	Trans Perf	Eng	Infra	R	L
91	RNP (Brazil)		Cons	5/27/20	8/20/20	Trans Perf	Eng	Infra	R	L
92	University of Central Florida		Cons	5/29/20		DMZ	Eng	Infra	E	L
94	NLM	TP	RA	6/5/20		Trans Perf	Eng	Bio	R	S
95	JGN-CCNY	TP	Cons	6/8/20		Trans Perf	Eng			
96	Yale		Cons	6/16/20	9/28/20	Arch, DMZ	Eng	Infra	E	L
97	University of Central Florida		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		Trans Perf	Eng/Proactive	Bio	R	L
101	LBL		RA	7/17/20		Perf	Eng	Infra	R	L
102	USDA		Cons	7/24/20		DMZ	Eng	Infra	R	L
103	VRO (LSST)		Cons	8/3/20	9/16/20	PS	O	Astro	O	S
104	Allen Institute/ Princeton		Cons	8/3/20		Perf	Eng	Infra / Bio	R	L
105	Lafayette College	KINBER	Cons	8/13/20		DMZ	Eng	Infra	E	S
106	FRGP	FRGP	Cons	8/21/20	9/28/20	DMZ	Eng	Infra	E	S/M
107	Arecibo/ UCF		Cons	8/21/20		Trans Perf / DMZ	Eng	Astro	R	L
108	KAUST		Cons	8/26/20		DTN	Eng	Infra	R	L

109	NIH	Globus	Cons	8/31/20		DMZ	Eng	Bio	R	L
110	FRGP	FRGP	Cons	9/2/20	9/8/20	Perf	Eng	Infra	R	L
111	Kent State University	OARnet	Cons	9/8/20		pS	Other (Student)	CS	R / E	S / M
112	U South Carolina		Cons	9/10/20		DTN	Eng	Infra	E	M
113	TACC	TACC, LEARN	Cons	9/10/20		Perf	Eng	Infra	R	L/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 In Progress Application Deep Dives

There is one ongoing Deep Dive report:

- **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 are still being drafted but 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.

7.B Related Activities

With seven completed Deep Dives to date, we now have a set of data regarding CI preparedness to help us scope future activities. This will include documenting:

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

Presentations were given to both CARCC [6] and the Quilt [9] on how the Deep Dive process works, and we plan to do another for NOAA in Quarter 3.

7.D Upcoming Deep Dives

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.

Prior to COVID-related travel restrictions, there were three 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 to take place during Spring, 2020.
- **University of Central Florida (UCF):** Staff from UCF approached EPOC to stage a Deep Dive for the campus.
- **University of South Dakota (USD):** Staff from USD approached EPOC to stage a Deep Dive for the campus.

Each of these are now tentatively scheduled for the second half of 2021, although an additional delay may be necessary.

Additional discussion surrounding to other events started, but due to the pandemic have been stalled indefinitely:

- **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.

EPOC has begun a number of conversations internally, with community members, and with our advisory board, to evaluate strategies to adjust our approach to Deep Dives in light of the current travel restrictions. We are considering experimenting with one-on-one Deep Dives, more limited in scope, via video conferencing. However, the Deep Dive process has always relied heavy on the synergies that occur when CI engineers and researchers are in the same room, so it is unclear how successful this approach would be, and the Advisory Board also emphasized their belief that this activity would lose much of its effectiveness without the in person component. In any case, we will continue to create additional training materials, including but not limited to video discussions, lecture materials, and templates, that can be adopted by campus or regional network staff to conduct Deep Dive interviews on their own.

7.E Metrics

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

Meet Date	Appl name	Public/Private	Audience	Offered or Req	Head Count	Issues Identified	Complete Date
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	Est. Y3Q3
2021*	Arizona State Univ, Sun Corridor	Priv	ASU, UofAZ, NAU, and Sun Corridor Network staff	Req			
2021*	Univ Central Florida	Priv	University Researchers & Staff, Florida Lambda Rail	Req			
2021*	Univ South Dakota	Priv	Staff from GPN, USD, SDSU, Black Hills State, and other guests	Req			
On hold	Oregon State Univ	Priv	TBD	Req			

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 Current Deployments

Different components of NetSage can be deployed in different ways, depending on the requirements of the customer. This quarter, the NetSage development team released NetSage 1.5.0, which was deployed on all of the EPOC partner deployments. NetSage 1.5.0 added the ability to look at Science Data based on known Science Projects, as defined in the NetSage Science Registry (<https://scienceregistry.netsage.global/>). It also included a Dashboard that enables viewing Top Talkers over longer time frames to better understand resource usage.

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

- **FRGP:** The FRGP flow data deployment of NetSage was made public in December 2019 at <https://frgp.netsage.global/>.
- **Great Plains Network:** 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. Members from the EPOC team met with GPN at the end of May to review these new features as well as to gather feedback about future work. This resulted in several new items tracked in NetSage development.
- **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>. We have begun to work with iLight to identify members on their network without ASNs so that we can properly tag them using Shared Whois Project (SWIP) data. This task was planned for completion in Quarter 2, but progress by iLight has been delayed.
- **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/>. Similar to iLight above, we have begun working with engineers to identify and properly tag members without ASNs on their network using Shared Whois Project (SWIP). This work will continue into Quarter 3.
- **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 asked again to delay additional conversations until Dec 2020.
- **OARnet:** OARnet's new CEO has asked to delay discussions of a NetSage deployment until 2021 as he deals with other priorities.
- **Southern Crossroads:** The SoX NetSage deployment for flow data started receiving data at the very end of June, and was made public in July at <https://sox.netsage.global/>. We expect to continue working with SoX to expand the datasets that we are collecting to include SNMP.
- **Sun Corridor:** A NetSage deployment is in discussion.
- **TACC:** Conversations with TACC about an EPOC deployment for flow data from network devices kicked off in April 2020. We already have been collecting Tstat data from their archive sites as a separate effort. Work on the flow collection began shortly afterwards on deploying the new containerized flow pipeline and this work was completed in early July, 2020. This data is now publicly 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.

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

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

8.B Network Performance Detection

In Quarter 2, EPOC staff continued to use 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. The current investigations we are pursuing are described as part of the Roadside Assistance Cases #90, 94, and 100, detailed in Section 6.B.

8.C Metrics

Table 3: Metrics for NetSage activities for Year 3 Quarter 2.

Where Regional	Data	Date Live	# Monitored Devices	# Large Flows	# Unique Src Orgs	# Unique Dest Orgs
GPN	SNMP, Flow	10/18	2 routers	6,314,873	2,412	1,693
iLight	Flow	4/19	5 routers	57,061,858	4,134	11,447
KINBER	Flow	11/19	2 routers	16,639,233	2,470	3,675
FRGP	Flow	1/20	1 router	37,624,891	3,555	5,912
SoX	Flow	7/20	3 routers	19,402,578	2,883	2,373
TACC (LEARN)	Flow	7/20	1 router	1,048,388	891	813
TACC (LEARN)	Tstat	1/19	4 head nodes	582,348	97	129
UHawaii Astro	Tstat	5/19	1 DTN	272,250	136	309
NCAR (FRGP)	Tstat	7/19	1 DTN	13,313,299	531	2,032
NERSC	Tstat	3/18	10 head nodes	957,049	113	116

9. Data Mobility Exhibition and Data Portal

In Quarter 1, we decided to formally add an Activity area to group the work on the Data Mobility Exhibition and the Data Portal. Ken Miller is the lead for this area.

9.A Data Mobility Exhibition

Data mobility is a critical component of the process of science. Being able to predictably and efficiently move scientific data between experimental sources, processing facilities, long term storage, and collaborators is a common use case that transcends the boundaries of research disciplines. The Data Mobility Exhibition(DME) is using reference data sets and existing campus CI components to highlight actual data speeds experienced by R&E institutions. Participants download, measure, and potentially improve their scientific data movement capabilities as part of this activity. 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/>.

The DME 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 means 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.

This effort was launched in September 2019 as part of the CC* PI Meeting/Quilt meeting. Approximately 206 tests had been performed during Quarter 2. EPOC is working with sites that request assistance to understand and improve data architecture, tuning, and usability of the resources for science use cases.

On the technical side, three of the either test 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. They can be found here: <https://github.com/vasv/dme-utils> These scripts provide a way to do automated testing via the cli instead of manual tests through the Globus GUI.

9.B Service Development - Modern Research Data Portal

Data portals provide a way for scientists to search for, discover, access, download, analyze, and publish scientific data. They are incredibly valuable for large collaborations, research groups, and indeed for entire fields. 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 exposes a set of scientific data. The EPOC Data Portal is based on the MRDP and uses Docker to containerize the functionality for easier deployment.

The current Pilot Portal uses Docker and supports a graphical front MRDP site with federated authentication and a custom pointer file to any existing site DTNs. This setup installs a custom project or site front end with an API call to existing and collaborative Globus Endpoints. Initial tests with the Pilot have been challenged due to end users unfamiliarity with Docker and the fact that many of the data sets to be accessed are not in a common DTN setup. EPOC staff are working with members of the Globus team to simplify our approach.

Initial investigation of this work started in early 2020 with a deployment at ESnet. This quarter, the University of Hawaii Astronomy group and LEARN member Baylor pilots are continuing, albeit slowly. Penn State may also be a potential pilot user to assist with the data sharing needed for their CryoEM system, however they are waiting for a more stable Pilot to work with.

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 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.

10.A Current Status of Managed Service Deployments

EPOC regional partners expressed 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 are on hold through the end of calendar year 2021. At that time, we will assess the status of how our partners would like to move forward in these areas.

The exception is the ongoing, although frequently interrupted, work to deploy perfSONAR nodes jointly with KINBER at Arcadia. KINBER is changing its business model, and we are in ongoing conversations to move this effort forward.

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

Responding to continued calls in the R&E community for perfSONAR training, EPOC staff members Southworth and Chevalier lead an introductory perfSONAR class for network engineers from the LEARN community on September 21-22, 2020. 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 also participated in hands-on exercises, made possible by utilizing the virtual lab environment hosted by the University of South Carolina.

11.B Border Gateway Protocol (BGP) Training

EPOC has received a number 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. EPOC solicited feedback on network engineer BGP experience via a survey sent out to the community late in this quarter. This survey will close in early Quarter 3 and help inform upcoming content and training opportunities.

11.C PerfSONAR

We continued our work with the Network Startup Research Center (NSRC) to create a series of instructional videos for perfSONAR, following on from their initial work in August 2016. EPOC staff are working collaboratively with both NSRC and the perfSONAR Documentation and Training team to update this material. During this quarter, members of the perfSONAR Documentation and Training team have volunteered to record sessions on various topics in need of refreshing. Once footage has been created it is sent to NSRC for processing.

11.D Additional Year 3 Training Plans

EPOC staff have started the internal work to adjust our training to adapt to these circumstances with a set of expanded training modules to meet community needs. The goals we have identified around the extended training include:

- Creation of new online content that can be shared beyond pandemic;
- Address known areas of confusion or popular areas that are not articulated well;
- Involve a large population (600+) of the CI engineering community through use of the CI-ENG mailing list;
- Work with regional and application partners to both advertise these sessions as well as understand needs for content; and
- Motivate a large population of R&E sites to adopt best community practices around CI technology approaches;

We envision creating a series of short (20 minutes to 1 hour) taped content pieces that can be held live but also shared via YouTube for asynchronous viewing. We are identifying content areas by analyzing our Consultation cases and the results of Deep Dives but will also be

discussing areas of interest with the community more broadly in future quarters. Our plan is that EPOC staff will coordinate this effort, but we might not always be the presenters. In fact, our preference will be to involve external speakers and voices whenever possible. And while we envision getting these activities up and running regularly while travel restrictions are in place, if they are successful, we would continue them indefinitely. This includes the creation of “evergreen talks”, talks that are still applicable over time, not dated (aka brown and withered) due to changes in technology. These will focus on routing policy, performance monitoring, or system design. EPOC continues to evaluate and curate possible talks in this space, particularly as we all will change our travel and training behaviors due to COVID-19.

One of the first of the new training videos will be on the topic of “How do I specify a 10G DTN”. This is meant to be a very pragmatic approach to the defining and purchasing of a data transfer node, with walk throughs and pointers to component selection decision points and reasoning.

12. Data Privacy and Security

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>.

Basic security measures are being maintained and there were no security incidents to report for Year 3 Quarter 2.

13. Reporting Against Deliverables

Table 4: Current deliverables and status updates for Year 3 Quarter 2.

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 (NEW)	C- FRGP (106)	Started Y3Q2, Compl Y3Q2
RA 3.2.7 (NEW)	C - FRGP (110)	Started Y3Q2, Compl Y3Q2
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, Ongoing
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 (NEW)	C - TACC (113)	Started Y3Q2
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, Ongoing
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 (NEW)	C- Kent State (111)	Started Y3Q2
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, Ongoing
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 3.6	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, Ongoing
RA 3.6.7	C- Penn State (87)	Started Y3Q1, Compl Y3Q2
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
	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 3.15	Sun Corridor RA	
RA 3.15.1	Sun Corridor (81)	Started Y3Q1, Ongoing
RA 3.15.2	ASU (83)	Started Y3Q1, Compl Y3Q2
RA 3.15.3 (NEW)	ASU (99)	Started Y3Q2, Compl Y3Q2
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	C - Allen Inst (74)	Started Y2Q4, Ongoing

RA 4.36	C - Reed (75)	Started Y2Q4, Ongoing
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, Ongoing
RA 4.46	NLM/ASGC (94)	Started Y3Q1, Ongoing
RA 4.47	CCNY/JGN (95)	Started Y3Q1, Ongoing
RA 4.48	Yale (96)	Started Y3Q1, Compl Y3Q2
RA 4.49	UCF (97)	Started Y3Q1, Compl Y3Q2
RA 4.50 (NEW)	126.com (98)	Started Y3Q2, Compl Y3Q2
RA 4.51 (NEW)	EBI (100)	Started Y3Q2, Ongoing
RA 4.52 (NEW)	LBL (101)	Started Y3Q2, Ongoing
RA 4.53 (NEW)	USDA (102)	Started Y3Q2, Ongoing
RA 4.54 (NEW)	VRO/LSST (103)	Started Y3Q2, Compl Y3Q2
RA 4.55 (NEW)	Allen Inst/Princeton (104)	Started Y3Q2, Ongoing
RA 4.56 (NEW)	Lafayette College (105)	Started Y3Q2, Ongoing
RA 4.57 (NEW)	Arecibo/UCF (107)	Started Y3Q2, Ongoing
RA 4.58 (NEW)	KAUST (108)	Started Y3Q2, Ongoing
RA 4.59 (NEW)	NIH (109)	Started Y3Q2, Ongoing
RA 4.60 (NEW)	U South Carolina (112)	Started Y3Q2, 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	Arizona State/Sub Corridor	Tentative 2021(COVID)
DD 2.2.2	Northern Arizona Univ	OBE
DD 2.3	LEARN Deep Dives	Ongoing
DD 2.3.1	Trinity University	Compl - Event Y2Q1, report Y2Q3
DD 2.3.2	Baylor	Ongoing - Event Y2Q4, report TBD
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	Tentative 2021 (COVID)
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.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	Tentative 2021 (COVID)
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 (NEW)	SNMP for TACC	Discussion ongoing
NS 1.3.5 (NEW)	Flow for TACC	Initial Y3Q2, Ongoing support
NS 1.4	NetSage for OARnet	Ongoing
NS 1.4.1	SNMP for OARnet	Discussion Year 3
NS 1.4.2	Flow for OARnet	Discussion Year 3
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 Y3Q2,, Ongoing support
NS 1.6	NetSage for KINBER	Ongoing
NS 1.6.1	SNMP for KINBER	Discussion Ongoing
NS 1.6.2	Flow for KINBER	Initial Y2Q3, Ongoing support
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
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/Data Portal	Note: Restructured Y3Q1
DP 1	Data Mobility Exhibition Setup	Ongoing
DP 2	Data Mobility Exhibition Support	Ongoing
DP 3	Data Portal	Ongoing
DP 3.1	Portal Prototype Development (Prev MS 1.5)	Development begun Y2

DP 4	Portal Deployments	Ongoing
DP 4.1	Portal with Baylor (LEARN) (Prev MS 2.3.3)	Started Y3Q1
DP 4.2	Portal with GPN member (Prev MS 2.5.2)	OBE
DP 4.3	U Hawaii Astronomy Data Portal MS	Started Y3Q1
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	Now DP 3.1	
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	Now DP 4.1	
MS 2.4	OARnet MS	TBD
MS 2.4.1	OARnet DT MS	On hold Year 3
MS 2.4.2	Testset Loan to Kent State	Expected Y3 (COVID)
MS 2.5	GPN MS	TBD
MS 2.5.1	GPN and KanREN DT MS	On hold Year 3
MS 2.5.2	Now DP 4.2	
MS 2.6	KINBER MS	Started Y1, Ongoing
MS 2.6.1	KINBER and Arcadia PS MS	Deployment and Training postponed (COVID)
MS 2.6.2	Testset Loan to Duquesne	Expected Y3 (COVID)
MS 2.7	Other MS Deployments	Ongoing
MS 2.7.1	Now DP 4.3	

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 until 2021 (COVID)
T 2.9	PS Training for LEARN	Compl Y3Q2 (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 (changed to 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	Ongoing
T 4.6	10G Easy DTN Video	Started Y3Q1, ongoing
T 5	Reworking Training during COVID	Ongoing
T 5.1	New plan development	Started Y2Q4, ongoing

14. Financials

Expenditures were below the predicted rates in Year 3 Quarter 2, primarily due to pandemic-related travel and hiring restrictions. We are in active discussions with our NSF Program Officer about this situation and plan to request a No-Cost Extension for the project later in Year 3.

Table 5: Expenditures for Year 3 Quarter 2.

Item	Univ	Jul-20	Aug-20	Sep-20	TOTAL
STAFF COSTS (INCLUDING BENEFITS, F&A)					
Schopf, Jennifer-PI	IU	3,007	3,007	3,007	9,021
Addleman, Hans	IU	12,186	12,186	12,186	36,558
Chevalier, Scott	IU	594	594	594	1,782
Southworth, Doug	IU	4,541	4,541	4,541	13,623
Moynihan, Ed	IU	1,890	1,890	1,890	5,670
Hubbard, Heather	IU	1,570	1,570	1,570	4,710
Doyle, Dan	IU	2,315	2,315	2,315	6,945
IU Dev Team	IU	8,132	8,132	8,132	24,396
IU OmniSOC support	IU	1,213	1,213	1,213	3,639
Zurawski, Jason	LBNL	13,539	10,888	9,450	33,877
Robb, George	LBNL	2,485	1,108	184	3,777
Miller, Ken	LBNL	8,696	7,309	6,503	22,508
TOTAL STAFFING		60,168	54,753	51,585	166,506