

EPOC Managed Services

Last updated September 13, 2019

ABOUT EPOC

Over the last decade, the scientific community has experienced an unprecedented shift in the way research is performed and how discoveries are made. Highly sophisticated experimental instruments are creating massive datasets for diverse scientific communities and hold the potential for new insights that will have long-lasting impacts on society. However, scientists cannot make effective use of this data if they are unable to move, store, and analyze it.

The Engagement and Performance Operations Center was established in 2018 as 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 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 five main activities

- *Roadside Assistance* 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 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;
- Provision of *managed services* via support through the IU GlobalINOC and our Regional Network Partners;
- *Coordinated Training* to ensure effective use of network tools and science support.

MANAGED SERVICES

EPOC is tasked with developing a set of service definitions for common research and education cyberinfrastructure components. The goal of these definitions is to provide guidance for our Regional Networking Partners to help us clearly gather requirements and design the service before we implement, maintain, operate (potentially for a fee), and upgrade the service as a benefit for downstream connectors. In doing so, the costs associated with design, specification, and installation can be ameliorated for a larger population that would otherwise be reticent to adopt 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 is targeting four use examples of Managed Services:

- perfSONAR: a widely-deployed active test and measurement infrastructure that is used by science networks and facilities around the world to monitor and ensure network performance.
- NetSage: a measurement and monitoring infrastructure that primarily uses passive information from SNMP and de-identified flow collectors to understand how circuits and archives are being used, and to debug performance issues.
- Science DMZ: Dedicated network infrastructure specifically configured for the security and performance implications required for scientific use cases.
- Data Transfer: PC-based Linux servers built with high-quality components and configured specifically for wide area data transfer.
- Network Testset: specialized hardware used to provision and validate network infrastructure.

There are several known barriers to adoption of these technologies:

- Initial Cost of Hardware: Hardware (servers, routers, switches, etc.) can be an expensive investment given typical drivers that consume resources. Cost is a factor for a medium to small institution that struggles to maintain basic IT infrastructure for mission critical applications.
- Expertise Gap: Knowledgeable staff that are capable of designing and operating infrastructure are often in short supply. The demands on these staff include, but are not limited to installation, operation, ongoing maintenance, data analysis, and troubleshooting both local and external problems
- Lack of Compelling Use Case: Smaller institutions may not have more than a handful of use cases that involve data-intensive scientific research across a wide area connection
- Inertia: The typical designs of campus are sufficient for 99% of the users, thus the addition of new technology for a single use case may not be a high priority.