

## Why a 100Gbps perfSONAR Testpoint won't run a 100Gbps Test

Contact Point: Brenna Meade

Audience: General

Last updated: August 17, 2021



### 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 six main activities:

- *Roadside Assistance and Consultations* via a coordinated Operations Center to resolve network performance problems with end-to-end data transfers;
- *Application Deep Dives* to work more closely with application communities and 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;
- *Data Transfer Testing/ Data Mobility Exhibition* to check transfer times against known good end points;
- *Provision of managed services* via support through the IU GlobalNOC and our Network Partners;
- *Coordinated Training* to ensure effective use of network tools and science support.

### Why a perfSONAR Testpoint won't run a 100Gbps Test

When planning a network architecture it is often tempting to build as big as possible with the mindset that maximizing the current options will aid in future proofing or hardware longevity. This approach is also seen as engineers design and deploy their testing infrastructure, for example, by deploying 100Gbps capable perfSONAR test points. If a site does deploy a 100Gbps perfSONAR testpoint, it is unlikely to achieve that performance in a production environment, as seen by the over 85 publically available 100Gbps perfSONAR testpoints in the perfSONAR Lookup Directory as of late 2021, which averaged throughput performance between 10 and 25Gbps. At this time, most end-to-end paths are not configured to support a 100Gbps flow, if

the path also supports production traffic then a 100Gbps test may have significant negative effects on production data movement. Therefore it is important to understand the path your

### **Issues When Running 100Gbps perfSONAR Tests**

When using perfSONAR testpoints in a production wide area network (WAN) environment, even if part of the end-to-end system is 100Gbps capable, we often see that the full path is not. For example, there may be a switch or router in the middle of the path that is only 10Gbps capable, or there may be a firewall or traffic shaper in the path that can negatively affect performance. Before considering the deployment of a 100Gbps perfSONAR testpoint, the capacity of the testing path should be verified.

However, even if the end-to-end path is 100Gbps capable, that much capacity may not be available at the time of the testing due to other traffic. When a 100Gbps perfSONAR testpoint runs a test it will fully saturate the link, which may cause interruptions to production services. Not only will the test not achieve 100Gbps, but the production traffic may experience additional latency or loss due to the test. Because of this, we do not recommend running 100Gbps tests on a link that has production traffic and is only 100Gbps.

### **EPOC Recommendation**

It is important when considering perfSONAR for your network that it is understood that perfSONAR is meant to be an example of the type of experience a user will have, and not a large scale network test such as those performed to validate a circuit. Though you may get 100G performance during a network capacity test, you will not see the same results from a perfSONAR test. Because of the issues currently seen with 100Gbps perfSONAR Testpoints, we recommend using one or more 10Gbps or 40Gbps Testpoints for more cost effective, practical, and accurate network performance testing. By deploying several smaller testpoints, more pieces of the network will be able to be tested and the tests themselves will more closely resemble what a science application is experiencing.

For those institutions working to future proof through investment in 100G or 200G technologies, the best investment is likely to purchase 100/200G network interface cards, instead of perfsonar nodes. In order to use these 100G or 200G interfaces to help maintain effective performance, they should optically limit the test server to at most 40G, again because we do not see accurate results in 100G testing. An added benefit of reducing the speed is that the tuning network performance becomes much easier and more effective.

For additional resources on improving performance testing please refer to the following resources

- <https://www.perfsonar.net/>
- <https://fasterdata.es.net/performance-testing/perfsonar/>
- <https://fasterdata.es.net/network-tuning/mtu-issues/>

- <https://fasterdata.es.net/network-tuning/tcp-issues-explained/>
- <https://fasterdata.es.net/science-dmz/>
- <https://epoc.global/wp-content/uploads/About-MTUs.pdf>