



Internet2 Performance Tools Case Study: LHC-Testing Challenges

www.internet2.edu/performance
www.ultralight.org


UltraLight is a collaboration of experimental physicists and network engineers working to provide the necessary network capabilities to enable petabyte-scale analysis of globally distributed data. Current Grid-based infrastructures provide massive computing and storage resources, but are currently limited by their treatment of the network as an external, passive, and largely unmanaged resource. UltraLight goals are to:

- Develop and deploy global services that promote the network as an actively managed component.
- Integrate and test UltraLight in Grid-based physics production and analysis systems currently under development in the Large Hadron Collider (LHC) ATLAS and CMS experiments.
- Engineer and operate a trans- and intercontinental optical network testbed, including high-speed data caches and computing clusters.

The Problem

High-energy physicists at the University of Michigan involved in the US-ATLAS experiment recently began monitoring the network infrastructure connecting their site and Brookhaven National Laboratory, the US-ATLAS Tier 1 location from which they receive their data. The physicists also began monitoring between their site and the other peer LHC institutions—called Tier 2 sites—with which the University of Michigan shares data. Each of the US-ATLAS Tier 1 and Tier 2 sites have deployed a pair of perfSONAR Performance Nodes for this purpose.

Using BWCTL, via perfSONAR-BUOY, regular monitoring tests were scheduled. Reviewing the results it was noticed that performance to every peer site was a mere fraction (1/10) of what it should be in the outbound direction. While it was now clear there was a performance problem, it was still necessary to pinpoint what caused the issue and then determine the best course of action to resolve it.



The Solution

A multi-domain troubleshooting team, which included staff from ESnet, Internet2, UltraLight, and the University of Michigan began the analysis task by using the NDT and NPAD tools to capture in-depth data that examined how TCP was operating over this path. These tools revealed a bottleneck in the path that was limiting the hosts' performance. By using servers and clients along the path, the divide-and-conquer methodology proposed in the Internet2 Network Performance Workshops, the team was able to quickly isolate the problem to a subset of links on the path.

The Result

Armed with this knowledge the Ultralight engineers began a detailed look into the switches and routers they controlled. "It was finally discovered that the forwarding engine on one of the switches had entered into a fault state causing all packets to be processor-switched," noted Shawn McKee. "The switch counters were not indicating loss, but the error condition had disabled the hardware line-rate packet forwarding feature of the switch and slowed outbound packets by a factor of 10. A simple reboot resolved the problem."

This type of 'soft failure' – not due to mis-configuration but an actual hardware problem – normally go undiagnosed and are the type of problems for which the Internet2 Performance Node is an ideal troubleshooting tool. "If a hardware fault causes work to stop, engineers fix the problem; if the network is limping along, engineers don't have time to jump in and fix it. In this case, the monitoring clearly showed an asymmetric behavior, which convinced the engineers that it was a 'real' problem – but the fault probably existed for weeks before it was finally caught and now resolved," says Rich Carlson, developer of the NDT tool and part of the troubleshooting team.

For more information about UltraLight, see www.ultralight.org; for more information on the Internet2 Performance Tools, see www.internet2.edu/performance.

Internet2 Performance Tools

The suite of Internet2 Performance Tools contains a user-friendly tool (NDT) that allows end-users to examine their local desktop/laptop computer's configuration and the local network infrastructure; with NDT, end-users generate trouble reports that 'snapshot' the problem as it is happening and eases the burden of system administrators and network operations center (NOC) staff by providing more diagnostic tools and procedures to help them quickly locate the source of real problems.

The Internet2 Performance Tools also include OWAMP and major parts of the perfSONAR-PS software suite. The perfSONAR-PS framework creates a measurement domain that seamlessly interacts with peer domain infrastructures, making the job of solving end-to-end performance problems on paths crossing several networks much easier to address. The perfSONAR-PS software package features: SNMP Measurement Archive, Lookup Service, and perfSONAR-BUOY with BWCTL.

Performance Toolkit

The perfSONAR Performance Toolkit is a downloadable, pre-configured collection of network performance tools. This software bundle, recommended and packaged by Internet2, was designed to collect, store, and analyze network performance. You can use this package to create a fully realized network performance node. When used as part of a periodic monitoring infrastructure, these tools can help identify changes in network performance or point out when upgrades are required. Instantaneous use can help NOC staff quickly resolve problems that end users discover and help put to rest the idea that the network is broken. For more information, see pmps.perfsonar.net/toolkit.



Internet2 Offices

1000 Oakbrook Drive
Suite 300
Ann Arbor MI 48104
T: 734.913.4250
F: 734.913.4255

1150 18th Street, NW
Suite 1020
Washington DC 20036
T: 202.872.9119
F: 202.872.6648