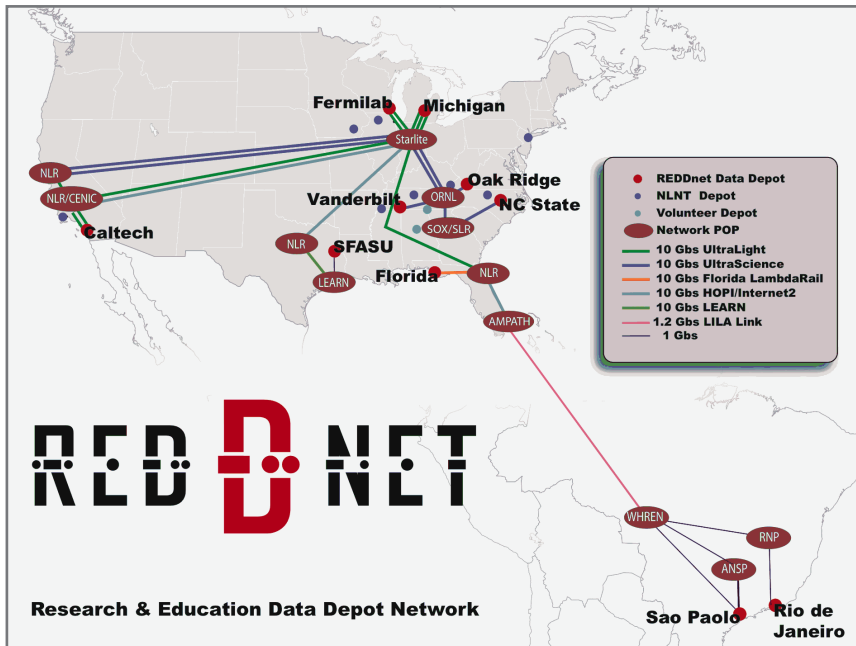


Internet2 Performance Tools Case Study: REDDnet's Distributed Storage Challenges



The solution

In collaboration with Internet2, REDDnet developers deployed portions of the perfSONAR Performance Node suite of performance and monitoring tools at several of its depot locations. Tools such as OWAMP, BWCTL, NDT and perfSONAR-PS can be used to diagnose existing problems as well as to create an infrastructure capable of long-term monitoring and analysis. Using these tools, the network path was examined from each depot's physical location to its nearest Internet2 Point of Presence (PoP). Regular testing using the perfSONAR-BUOY BWCTL throughput monitoring service revealed that baseline performance failed to meet expectations at many facilities, generally getting 50–100 Mbps versus the expected 1 Gbps rates. Digging deeper, using tools such as NDT, OWAMP and Traceroute, the team was able to pinpoint and resolve several network problems ranging from asymmetric routing to failed network devices, each of which contributed to the overall poor network performance.

Using the methodology and tools proposed in the Internet2 Network Performance Workshops, the group started the debugging process:

- 1 Establish baselines using regular BWCTL tests.
- 2 After a couple of days to get baselines, start using NDT to test into a “known to be good” area (e.g., the Internet2 backbone). Each REDDnet depot tested to the closest Internet2 PoP and the results were compared.
- 3 Divide and conquer: (a) Test into the Regional network. (b) Test into the campus network. (c) Test into the building network until the performance starts to resemble what it should be for a 1Gbps network.

Research and Education Data Depot network (REDDnet) is a National Science Foundation-funded infrastructure project designed to provide a large distributed storage facility for data-intensive collaboration among the nation's researchers and educators in a wide variety of application areas. Its mission is to provide “working storage” to help manage the logistics of moving and staging large amounts of data in the wide area network. The system supports researchers who need to transfer data inter-institutionally to their remote collaborators or share large data sets with their collaborators for limited periods of time (ranging from a few hours to a few months) while they work on that data set.

The problem

REDDnet provides distributed storage capabilities called *depots* at key network hubs, including many U.S. institutions participating in the Large Hadron Collider (LHC) project. Scientists at these sites, involved in either the LHC CMS or ATLAS experiments, use the REDDnet depots to access their LHC data. For months, Paul Sheldon, a physicist with Vanderbilt University and Principal Investigator for REDDnet, was frustrated by the lack of optimal network performance when accessing his data. Testing showed that his connection was only achieving 100–200 megabits per second (Mbps)—a fraction of the 1 gigabit per second (Gbps) capable machines and connectivity Professor Sheldon's team had invested in.

The result

The performance tools work by breaking long network paths into smaller, distinct domain sections. In doing so, the tools helped REDDnet more accurately isolate and address issues—both locally, and at the campus and regional level. After testing in this manner with a couple of locations (note: this is an ongoing exercise; examples are given only from Vanderbilt), the team was able to determine:

- Performance leaving Vanderbilt was vastly different than performance entering Vanderbilt, differing by about 50–100 Mbps.
- Both directions were achieving only 100–200 Mbps instead of close to 1 Gbps.
- Testing from Vanderbilt to any other REDDnet location showed this poor behavior, as did testing from Vanderbilt to the Internet2 PoP in Atlanta and from Vanderbilt to the Southern Crossroads (SOX) GigaPoP.
- Testing from Vanderbilt to other locations *within* Vanderbilt **did not** show this poor performance!

This suggested the problem was isolated somewhere between the Vanderbilt border and the SOX uplink to Internet2. In working with REDDnet and SOX network staff, problem links were identified and corrected, which allowed performance to begin to approach normal levels. Further debugging suggested that hosts, routers and switches in the Vanderbilt path should be examined to be sure they were configured properly. After increasing the size of some buffers, network performance for the In and Out paths matched very closely. Re-testing to SOX, Internet2, and eventually other REDDnet depots, showed performance had gone from 100–200 Mbps to over 600 Mbps to select locations. Testing is still not complete, but this proved the usefulness of the Internet2 Performance Node for isolating problems.

As a result of this work, Professor Sheldon noted, “The REDDnet project is deploying storage depots on national and international high performance networks to support important scientific research. We discovered that it is important to *constantly* monitor and debug the performance of the network connections to and from our depots, and tools such as perfSONAR and NDT have been extremely valuable in this effort. We made little headway on this problem until we deployed these tools.”

FOR MORE INFORMATION

REDDnet: reddnet.org

Internet2 Performance Tools:

internet2.edu/performance

perfSONAR Network Performance

Toolkit: [internet2.edu/performance/](http://internet2.edu/performance/pS-NPToolkit)

[pS-NPToolkit](http://internet2.edu/performance/pS-NPToolkit)

Internet2 Performance Tools

Internet2 Performance Tools make it easy for everyone to participate in the debugging process. The tools can be used to isolate faulty components, identify poorly configured hosts or allow authorized users to examine performance data from—and between—peer networks.

The suite contains a user-friendly **Network Diagnostic Tool** (NDT) that allows end users to:

- Examine their local computer’s configuration and local network infrastructure.
- Generate a “snapshot” of the problem as it happens, providing more diagnostic tools to help locate the source of the problem more quickly.

Two traditional NOC-style tools, **BWCTL** and **OWAMP**, make it easy to examine WAN paths. Tests can be run to BWCTL and OWAMP servers inside the Internet2 Network (and other backbones, such as ESnet). Remote servers, located at other points in the network, may also be accessible. These tools archive test data for later analysis.

Also included are important pieces of the **perfSONAR-PS** software suite. This framework enables the creation of *measurement domains* that seamlessly interact with peer domain infrastructures. This makes solving end-to-end performance problems on paths crossing several networks much easier to address. The package features: SNMP Measurement Archive, Lookup Service, and perfSONAR-BUOY.

The **perfSONAR Network Performance Toolkit** (pS-NPToolkit) is a downloadable, pre-configured collection of network performance tools. Recommended and packaged by Internet2, the bundle is designed to collect, store and analyze network performance. The package can be used to create a fully realized **perfSONAR Performance Node**. When used as part of a periodic monitoring infrastructure, the tools can identify changes in 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.