

perfSONAR

Performance-Monitoring Middleware

PerfSONAR Helps Accelerate Big Science Collaborations

In the arena of high-performance networking, it's easy to track down "hard failures," such as when someone breaks or cuts through a fiber link. But identifying "soft failures," like dirty fibers or router processor overload, is challenging. Such soft failures still allow network packets to get through, but can cause a network to run 10 times slower than it should. They also account for the majority of performance issues that users experience.

A global collaboration consortium led by the U.S. Department of Energy's Energy Sciences Network (ESnet), GÉANT2, Internet2 and Rede Nacional De Ensino e Pesquisa (RNP) has now developed a network performance monitoring and diagnostic system called perfSONAR that is helping network engineers identify bottlenecks, allowing them to make relatively small tweaks to gain significant speedups.

U.S. perfSONAR development has been advanced via partnerships between the University of Delaware, ESnet, Fermi National Accelerator Laboratory, Internet2, and the SLAC Linear Accelerator Laboratory. Developed with usability in mind, a perfSONAR Performance Node boots from a CD, uses a low-cost Linux computer as a host and takes just 10 minutes to configure.

"Once it's up and running, perfSONAR can perform regular tests of a network," said Brian Tierney, an ESnet computer scientist, who is based at the Lawrence Berkeley National Laboratory. "Basically every time we have worked with someone to set up perfSONAR and run some bandwidth tests, they have found what I call a 'soft failure,' where bandwidth on some path is three to 10 times slower than expected."

Tierney has been developing tools to assess network performance for more than 10 years. These ongoing tests help differentiate temporary glitches from ongoing configuration problems. He notes that oftentimes soft failures are not obvious and can only be detected with close inspection.

Among the types of problems found so far at various universities and national laboratories around the U.S. are:

- Multiple cases of bad fibers
- Port-forwarding filter overloading a router and causing packet drops
- Under-powered firewalls which could not handle the amount of incoming traffic
- Router output buffer tuning issues
- Previously unnoticed asymmetric routing causing poor performance
- Under-powered host (doubled performance by switching to jumbo frames)

One of the largest upcoming networking challenges for the high energy physics community is transferring and accessing large datasets related to experiments at the Large Hadron Collider (LHC) at CERN in Switzerland. Once the LHC goes into full production in late 2009, terabytes of data will flow from CERN to Brookhaven National Laboratory (BNL) in New York and Fermi National Accelerator Laboratory in Illinois—both Tier 1 U.S. LHC sites.

From Europe to the U.S. Tier 1 sites, the data will traverse two networks, USLHCnet and ESnet. The data will then be sent to five other centers, known as Tier 2 sites, in the U.S., from which physicists around the nation will be able to access and study the data. From Tier 1 to Tier 2 sites, LHC data will traverse the ESnet and Internet2 backbones and various regional and local area campus networks. For example, LHC data from Fermilab destined for the University of Nebraska—Lincoln (UNL) will traverse Internet2, ESnet, Great Plains Network and the UNL campus networks to reach researchers.

"If we don't perform well, it slows everybody down. Physicists want the data to arrive as fast as humanly possible, if not faster," said Shawn McKee, a high-energy physicist who is also director of the ATLAS Great Lakes Tier 2 Center at the University of Michigan.

In preparation for moving and analyzing LHC data, the U.S. ATLAS Project is simulating what happens inside the detector on supercomputers and moving this information across multiple networks to ensure that everything is working properly. Among the millions or billions of particle collisions, a handful will be "unusual events," or extremely rare phenomena which will provide key insights into the origins of matter in our universe.

According to McKee, data flowed into the Michigan center at 900 Mbps, but tests on outgoing data showed rates of only 80–90 Mbps. Using the perfSONAR measurement hosts along the path, McKee and his colleagues were able to eliminate potential sources of trouble. Regular tests of the BNL-to-Chicago path via ESnet and Internet2 showed no problems, and the internal BNL path also appeared to be performing at speed. However, perfSONAR tests showed that something was wrong with the segment between Chicago and the Michigan site. Because all of the centers were running the same perfSONAR software, they were able to easily compare data.

"We had the impression that we had a problem, that the data was not moving out as fast as it was moving in, but we couldn't find out why. It was really unusual," said McKee, who notes that his team initially thought the larger networks were dropping packets. However, "counters" on the network routers showed that all the data was going through—just at one-tenth the expected rate.

"Finally, we thought, 'It's not the network—it must be us'," McKee said.

The Michigan team finally narrowed the source of the problem down to a fault in hardware forwarding on one of the 10 Gbps blade servers. Too many routes had been loaded onto the server so instead of forwarding the data, it sent each stream to a processor, which then made a software decision about each transmission. This slowed the entire process. Although the server was generating bug reports, there were no error messages indicating the problem. With help from colleagues at Caltech, McKee's team found the problem and the fix.

While this is a successful example of perfSONAR's capabilities, it also highlights one of the limitations: Although the system can find that a problem exists, it is not as good at pinpointing the exact cause of the problem. But this situation will improve, Tierney said, as more perfSONAR measurement points are installed on various networks.

"With perfSONAR, we can create a persistent baseline of performance in all segments of the network and see if any

changes arise," McKee said. "We can look at the ends of the network and, if there is a problem, run on-demand tests using perfSONAR on the suspect segment."

By adding more perfSONAR nodes, network engineers can divide the paths into smaller and smaller segments, helping to narrow down exact problems locations.

In addition to detecting soft failures in networks transporting data from the LHC, perfSONAR has also been used to identify bottlenecks in networks connecting the upcoming Daya Bay Neutrino Experiment, in Southern China, to computing and mass storage systems at DOE's National Energy Research Scientific Computing Center (NERSC) in Oakland, California, where data from the experiment will be analyzed and archived. Neutrinos, subatomic particles that widely populate the universe, are currently a mystery to scientists. The Daya Bay experiment will help researchers learn more about the properties of this puzzling particle, as well as its role in the universe.

"Before perfSONAR, it usually took several days to pinpoint the source of a bottleneck when massive datasets were transferred across multiple networks. We had to work with operators of each network to identify the problem and have it fixed," says Jason Lee, a NERSC network engineer. "Because perfSONAR actively and automatically searches for problems, we can quickly find choke points and immediately know who to contact to get it fixed."

"At first, I was kind of amazed at the number of soft failures we found using perfSONAR, but then I realized this is exactly what we were hoping to be able to do when we first started talking about perfSONAR 10 years ago," Tierney said. "Of course, in a way this makes our jobs harder, as perfSONAR finds more problems for us to fix."

ESnet interconnects more than 40 DOE research facilities and dozens of universities across the United States, also providing network connections to research networks, experimental facilities and research institutions around the globe. GÉANT2 and RNP are the research and education networks of Europe and Brazil, respectively. Internet2 is an advanced networking consortium that operates an advanced nationwide network connecting the U.S. research and education community.

NERSC is the DOE's flagship supercomputing center for unclassified computational science and NCCS is a DOE Leadership Computing Facility. Data collected from the Daya Bay experiment will also traverse ESnet on its way to NERSC.