



www.internet2.edu

When 99% Isn't Quite Enough:

A Case Study for E2Epi

On October 2, 2002, Internet2 began netcasting the EDUCAUSE 2002 conference from Atlanta; Internet2's President and CEO, Doug, was scheduled to give a talk (The Development and Future of the IT World: Higher Education's Role) on October 3 that would be presented from Internet2's Ann Arbor offices. Jon and Chris were the Internet2 staff responsible for netcasting the event.

During the test phase before the netcasting began, Chris told Jon that the netcast transmission quality was poor; Jon's video transmission equipment showed significant packet loss but, when he tried to pin down the location of the network problem, he had no luck. Ping was showing no packet loss and Iperf tests showed only 1% loss.

Not knowing what else he could do, Jon sent out an S.O.S. to the **E2Eperf Interest Group** on October 3:

"We currently have a network problem that is proving hard to pin down. Ping (every .1 sec, 1300 byte packets) shows no packet loss. From 199.77.176.251 to Ann Arbor (say 207.75.164.56), there is some problem. Iperf showed only 1% loss. Our video transmission equipment shows major packet loss. Any thoughts? My guess is some type of traffic shaping or perhaps back to back packets problem."

Russ and Matt, members of Internet2's End-to-End Performance Initiative (E2Epi) team, received the request for information; they called Jon to ask for details on the problem – what tests had he already run? What were the parameters? What were the results? What were the IP addresses and applications involved?

They began looking at the path: Russ ran a **jalaam test** that indicated some packet loss between Atlanta and Ann Arbor but it only

gave an indication of where the loss was occurring. Matt had an Iperf server available at the SoX (Southern Crossroads) gigaPoP, which serves Atlanta. Matt began running Iperf tests; he eventually discovered that 1% loss occurred from Georgia Tech to Ann Arbor but not the other way around. Matt also discovered that the loss was 1%, regardless of applied load – the loss remained constant between 5 and 50 mbps.

After discovering that the loss was occurring on the Atlanta to Ann Arbor route, Matt tested various parts of the path. First, he checked the Abilene Weather Map, which displays loss/error status on internal links. Then, he checked the connection from Abilene to Ann Arbor (via the Merit gigaPoP); no problems on the Ann Arbor end. Matt verified loss from Indianapolis to the SoX gigaPoP.

At this point, Matt and Russ began researching the SoX gigaPoP end of the path. They looked at the Abilene Atlanta router interfaces and found that the interface to the SoX gigaPoP was showing about 1% errors in the inbound direction only since the counters were last reset some time before – the interface was showing CRC (cyclic redundancy checks) errors. The router performs a checksum to ensure the packets are good; in this case, it was throwing away 1 % of the packets as corrupt.

Matt spoke with the Abilene NOC staff and found that this was a known problem with some optical gear through this gigaPoP; the SoX gigaPoP had been working on the problem. Matt spoke with one of the SoX gigaPoP network staff, who indicated that there had not been enough use of real-time applications on that link that need zero-loss performance to justify the time and expense of fixing the problem in a hurry.

What is the E2Eperf Interest Group?

The E2Eperf Interest Group is an email list for those interested in end-to-end issues. It was started to help form Internet2's End-to-End Performance Initiative (E2Epi) and continues today as a place for discussions of end-to-end performance problems. To subscribe to the mailing list:

- Send a message to the address: listproc@internet2.edu
- Do not include anything in the subject line
- In the body of the message include the line: `subscribe e2eperf-interest FirstName LastName [i.e., subscribe e2eperf-interest John Smith]`

Once you are subscribed, you may send a message to the list by using this address: E2EPERF-INTEREST@Internet2.edu

Archives of the E2EPERF-INTEREST list can be found at: <http://archives.internet2.edu/>

What is a jalaam test? Jalaam is a company that has end-to-end network testers, called Apparenet, that try to determine the performance and problems of each node along the path from the test location to a destination. Jalaam is working with the Internet2 community to test and improve their product.

Why is 1% loss significant? A video feed does not need much speed, only from 128K to 3 Mbps, but *any amount* of loss causes problems. H.323, the application that Doug was using to give his presentation, is an example of an application family that does not tolerate loss. Loss is significant to any real-time application, except voice. This affects remote labs, video conferencing, and anyone trying to do high bandwidth over long distance (because of current TCP characteristics). There are several new high-speed TCP algorithms under development that may fix this problem but it will be at least 3-5 years before they are in common use. The concept of zero-packet loss tolerance by some application families needs to be communicated more clearly to a broad audience.

While all the testing was going on, Jon and Chris managed the netcast of the conference. Presentations were transmitted live; the packet loss caused major disruption in the quality and effectiveness of the netcast. Sections of presentations lost both audio and video signals at times; as such, the netcast could be deemed a failure.

By 5:00 p.m. on October 3, Russ contacted Jon to inform him that the problem was known and unfixable: the SoX gigaPoP had a problem with their fiber connection and fixing the problem would be costly and time-consuming. The SoX gigaPoP did not communicate to outside world that the problem existed – there is no method by which they could have done so. If the conference planners (especially their networking/video engineers) had known of the problem, they could have tried to reroute the traffic, changed the conference site, or warned users that H.323 (and any other real-time) applications were too sensitive for the network.

Recommendations

In the discussion that ensued after Jon posted the problem to the E2Eperf Interest Group, several recommendations were made and issues raised:

- While developers/users of applications want the network to run perfectly – high speed and zero-loss – network engineers feel that application designers need to realize that the network never will be perfect; applications need to be more robust to withstand common network errors.
- Applications users, on the other hand, want network engineers to understand that there are families of applications that require near-zero packet loss but are not speed-driven (so systematic loss is significant).
- Users and network engineers must keep each other informed of changes if the network is to run smoothly. Engineers should think of how a modification or system failure could affect users and inform them of the change so that they can make adjustments, where possible;

users should inform network staff when they plan to use a new type of application to ensure that the network is capable of supporting it.

- Communication is important but education is essential – both users and network engineers need to understand that different types of applications have different network needs; users should be able to tell network engineers the type of application they will be running and have that convey information about the necessary speed and packet loss-resistance of the application family so a usable path is available.
- In response to the requests from several members of the E2Eperf Interest Group after this problem occurred, more links from the gigaPoPs have been added to the Abilene Weather Map; users need to double click on the node to see activity at the gigaPoPs that connect to that node.
- Members of the E2Eperf Interest Group felt that campuses and gigaPoPs should create a posting location to alert people to known or long-standing problems.
- Knowing the tolerance for an application is important. (How can users find out what the tolerance is? Is there a test one can run or is it all trial and error? Once one has discovered the tolerance, how can users share knowledge with others? What does the application require to work well and at what point does it begin to degrade or fail?) There has been discussion of creating information on the network requirements and tolerances of specific types of applications.
- In solving this problem, Russ and Matt had to look at the router proxy; users can issue commands to show data on interface statistics and get that information back. Matt and Russ were able to get the Atlanta Abilene router to show them the interface statistics and that was where they initially pinpointed the 1% loss. As a result of this problem (and its solution), this information is now available on the Weather Map.
- Keep track of loss/error statistics; Matt suspects that Jon's ping tests didn't show loss because it wasn't running long enough to catch the 1/100 loss.

Interested in learning more about E2Epi?
See our website at:

<http://e2epi.internet2.edu/>

There you'll find links to recent projects (like **E2E piPEs**) and presentations, as well as other information related to end-to-end performance issues, measurement tools, upcoming events, and related activities.