Correlating Internet Performance Changes and Route Changes to Assist in Trouble-shooting from an End-user Perspective

Les Cottrell, Connie Logg, Jiri Navratil, Stanford Linear Accelerator Center

With the growth of world wide data intensive scientific collaborations, there is a need to transfer large amounts of data to and from data repositories and collaborator sites around the world. To effectively enable such transfers, high speed, efficient, predictable networks are needed. In turn, these require excellent performance monitoring to ensure continuous optimal network performance for the applications to run. One tool/infrastructure that has been successfully used for such monitoring of critical paths (i.e. paths to sites that an organization requires excellent network performance with) is IEPM-BW [1]. Based on experience gained from the achievable throughput monitoring in IEPM-BW, we developed ABwE [2], a tool to enable quick (< 1 second), low impact (40 packets) measurement of available bandwidth. Using ABwE we have been able to quickly (within minutes) visually identify significant changes in available bandwidth on production links with up to 1Gbps bottlenecks. Investigating such changes, in particular degradations, we have found, not surprisingly, that many can be associated with route changes. Once such a significant performance change is discovered, the main problem for the end-user (e.g. network administrator at an end-site) is to: gather relevant information to identify the magnitude of the change; the time(s) it occurred; record the before and after routes; see if the change affects multiple paths; discover common points of change in the paths; identify the probable relevant Internet Service Providers; and report this information to the appropriate Network Operations Centers (NOCs). In our experience once the above has been done, the NOCs are fairly quick in responding with the cause of the change and often a fix. We have therefore developed a set of tools to facilitate the above process. The tools measure traceroutes at regular intervals, record them in an archive, and provide tools to enable simple visualization, navigation and integration with other tools such as ABwE and IEPM-BW, and a topology display. This presentation will present this set of tools and discuss their effectiveness together with examples of their utilization.

Methodology:

A monitoring host (with a 1 GE interface) has been set up at SLAC. Various remote hosts (40-50) at collaborating sites around the world have been chosen to perform tests to. Accounts with ssh access have been set up at these remote hosts. At regular intervals Ping, TCP transfers (using iperf), file transfer tools such as bbftp, and available bandwidth estimation tests (ABwE) are run. In addition, forward AND reverse trace-routes are run approximately every 10-12 minutes between SLAC and all the remote hosts. The results of these tests and the trace-routes are analyzed to: identify unique routes and assign route numbers to them for each remote host; identify significant route changes; and turn the data into more useful formats (web browsable for users, text format to embed in email to Internet Service Providers, log format for debugging etc.). This data is stored with the measurement time in an archive.

Visualization:

The simplest visualization technique involves time series graphs of the ping minimum and average Round Trip Times (RTTs), the results of the achievable and available bandwidth test results, the file transfer throughputs, and indicators on the time series graphs which indicate when the route to or from a node has changed. This allows for visual correlation of significant

changes in RTT, available and achievable bandwidth, and route changes in one or both directions.

A traceroute change to one node is often reflected in the traceroutes to other hosts. We therefore create a daily summary table that has the hour of day in columns and the remote hosts in the rows. Each entry (there can be multiple for each box representing an hour) provides a dot to identify that the route has not changed from the previous measurement or the new unique route number if the route has changed. The first measurement for each day is displayed with its route number. This very compact format enables one to visually identify if several routes changed at similar times, i.e. route numbers appear in one or two columns for multiple hosts (rows).

To facilitate further investigation of changes, in this table, there are highlighted links that allow one to: select the forward or reverse routes; view all the traceroutes for a selected remote host (as a color coded web table); access text suitable for attaching to trouble reports; review the log files; review the route numbers seen for a given host together with when last seen; view the available bandwidth time-series for the last 48 hours; navigate back through the archives for reports for previous days (the archives go back many months); and to select times and remote hosts for which one wishes to view topology maps for the selected routes and times. The topology maps display the hop routers colored by Autonomous System, provide the router and end host names by "mouseover", and provide the ability to zoom in on more complex routes.

Utilization:

This set of tools has been in production use at SLAC for several months. It has already been successfully used in several problem incidents and is being enhanced as a consequence of its use. We will report on specific examples illustrating how the tools have been used to identify and pin-point performance problems in various networks. In some of these cases the problem went unidentified for several days, but once identified and reported, the problem was fixed in hours.

Future plans:

Work is in progress to automate the identification of significant changes, and to automatically assist in gathering the associated relevant information (e.g. current, and trace routes before and after a performance change, time and magnitude of the change, topology map, and time series plots of the performance changes). This will be gathered into email and a web page and sent to the local network administrator. We expect to report on progress with this automation by the time of the conference. We are also analyzing the ratio of significant performance problems caused by route changes and vice-versa, and the duration of significant performance degradations, and will also report on this.

References:

[1] *Experiences and Results from a New High Performance Network and Application Monitoring Toolkit*, Les Cottrell, Connie Logg, I-Heng Mei, SLAC-PUB-9641, published at PAM 2003, April 2003.

[2] *ABWE: A Practical Approach to Available Bandwidth Estimation*, Jiri Navratil, Les Cottrell, SLAC-PUB-9622, published at PAM 2003.