PingER History and Methodology

R. Les Cottrell, Connie Logg and Jerrod Williams
Stanford Linear Accelerator Center, 2575 Sand Hill Road, Menlo Park, CA 94025, USA
cottrell@slac.stanford.edu, cal@slac.stanford.edu, and jerrodw@slac.stanford.edu

Abstract

This paper describes the methodology of the PingER toolkit/project. It provides some examples of how the results have been used to identify evolutionary changes in WAN response times and connectivity quality over the past 10 years, as well as some of the challenges faced in maintaining and deploying the system.

Introduction

In the early 1990s with the growing utilization of the Internet for communications among High Energy Physics (HEP) collaborators, SLAC started the development of a system for evaluating the quality of Internet connectivity between SLAC and other HEP Institutions.

Around 1996, SLAC's HEP collaborators became interested in the system. Several of the collaborators got together, proposed the PingER project [1] and obtained funding for it. PingER thus became the first Internet End to end Performance Monitoring (IEPM) tool available to the HEP community. Subsequently others using the Internet around the world became interested in it, and its deployment soared to encompass most of the regions of the world.

Contributing to PingER's popularity is the fact that the monitoring toolkit is very simple in concept and uses the well known, ubiquitous Internet ping facility. It requires no software or special hardware for the remote hosts to install, configure, or maintain, and, it generates a low level of network traffic thus minimizing the impact on low performance links.

Today the PingER project includes hosts at 35 monitoring sites in 12 countries that are monitoring hosts in 80 countries. These countries between them comprise 75% of the world's population and over 99% of the world's Internet connected population.

Methodology

Each of the monitoring sites chooses a set of Remote-hosts of interest to them. In addition a set of representative hosts (Beacons) have been chosen for various regions. These Beacon hosts (~ 70) are monitored by all Monitoring-hosts, thus providing measures of world-wide performance with respect to the Beacon hosts. In total there are over 800 Remote-hosts monitored at almost 500 sites in 80 countries.

Figure 1: Graphical representation of the PingER architecture
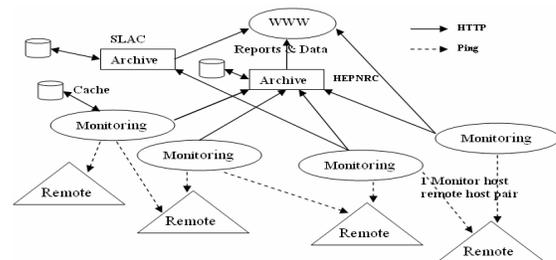
PingER's architecture (Fig. 1) defines 3 types of hosts that interact to conduct the monitoring. These hosts are the Monitoring-hosts, the Remote-hosts that are being monitored, and the Archive/Analysis hosts. By default, every 30 minutes each Monitoring-host sends eleven 100Byte pings and ten 1000Byte pings to each of the Remote-hosts being monitored by it. The 1000Byte pings are not usually sent to Remote-hosts with poor connectivity, thus reducing the impact on their networks. The first ping is used to prime name server caches and is then discarded. The Round Trip Times (RTT), losses, out of order and duplicate packets from the pings are recorded locally in flat files on the Monitoring-hosts. The data is gathered from the Monitoring-hosts, on roughly a daily basis, by the Archiving-hosts at SLAC and FNAL.

Figure 2: Example of the options menu from the www.slac.stanford.edu/cgi-wrap/pingtable.pl visualization tool. This shows the ability to select the specific metric, whether to give results for individual remote hosts or aggregate hosts by site, select the packet-size, choose the set of Monitoring-hosts and Remote-hosts, the time window for averaging data points, etc.


Figure 3: The monthly average RTTs (in ms) reported by pingtable.pl. The values are colored to assist in identifying problems. Column headings are clickable allowing the user to sort the data by the selected column.

Monitoring-Site	Remote-Site	Sep2003	Aug2003	Jul2003
SLAC	DESY	166.344	162.983	160.022
SLAC	MAGIC.MN	858.872	836.933	828.841
SLAC	SAVBA.SK	170.113	167.693	180.759
SLAC	UCSC	5.138	32.883	20.648
SLAC	UIB-NO	192.603	189.932	197.806
SLAC	MITLNS-MIT	80.570	82.636	82.232

The Archive-hosts also provide tools for the interactive analysis and web presentation of the data. These include navigational drill down displays, sorting tools, and the ability to select data by various aggregations and time-spans. For example, the hosts can be aggregated into groups by country, region, and/or affinity groups (e.g. Silk Road project [2] countries, sites collaborating with a HEP experiment, etc.). The user can select the data to be viewed by various characterizations (e.g. metric, time window and granularity) (see Fig. 2). The data can be viewed in tabular (see Fig. 3) and graphical formats, and also downloaded in tab separated value (.tsv) form for further analysis with tools such as Excel.
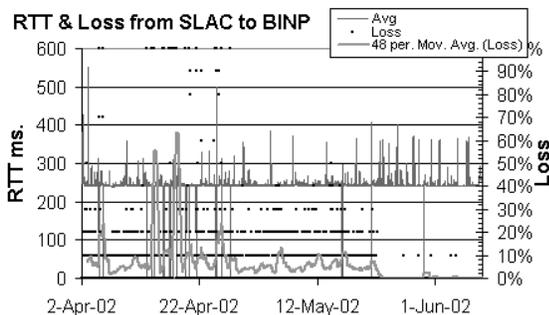
Figure 4: Network performance between ESnet Monitoring-hosts in the U.S. and Remote-hosts in the U.K
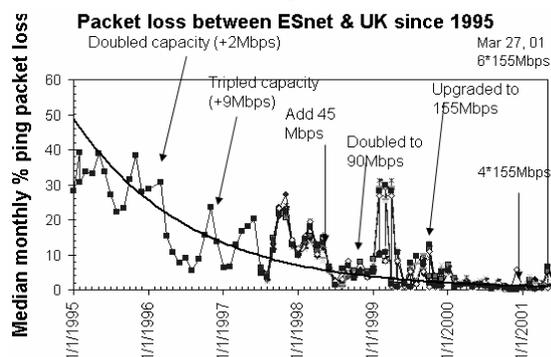
Uses and Examples

Over the past six years, the information gathered by the PingER project has been used in several ways. For example, it has been used to track network infrastructure changes. Packet loss measurements made from Energy Sciences network (ESnet) Monitoring-hosts to Remote-hosts in the U.K (Fig. 4) were able to chronologically pinpoint moments of network upgrades and effects of the upgrades (an increase in bandwidth of a factor of 300 dropped the packet loss from 35% to < 1% in 6 years), as well as the reduction in congestion during holiday seasons

Figure 5: Effects of network upgrades made to the BINP network in May of 2002. The thick light shaded line is to guide the eye and is the rolling average loss for the previous 24 hours

(as observed by the reduction in average packet loss). Further, the observed, close tracking of the losses on multiple paths between the two regions also illustrates the effectiveness of choosing Beacon hosts to represent a region.

PingER data has also been used to illustrate the need for upgrades to a network. The KEK/Tokyo - BINP/Novosibirsk link was experiencing heavy packet loss of around 15%. Based on a presentation of those findings, a successful recommendation was made to the policy/funding people to increase the bandwidth from 128Kbps to 512Kbps in an attempt to reduce the loss. In May of 2002, the upgrades were made to the link, and the packet loss was reduced to roughly 0.1% loss (see Fig. 5).

PingER has also been used to illustrate the difference in performance between developed and developing countries, sometimes referred to as the Digital Divide [3]. A by-product of this process has been the discovery that not all sites that are located in developing regions are seeing the negative effects of the Digital Divide. For example, the Middle East developing region includes sites being monitored in Israel which have better Internet connectivity than other sites in that region. Anomalies such as these can be used to justify the need to fund projects that increase the quality of internet connectivity to developing regions. If one site in a certain developing region can attain credible connectivity then other sites in that region should be able to have better connectivity as well.

Other practical uses of PingER include selecting Internet Service Providers (ISPs) for residential cable/DSL services and monitoring the accessibility to and changes made to ISPs from central hosts. As a troubleshooting tool, PingER has been used to discern if a reported problem is network related, identify the time the problem started, whether it is still occurring, and provide quantitative analysis for ISPs. It can also identify common patterns such as step functions, periodic network behavior, and recognize problems that synchronously affect multiple sites.

Challenges

One major challenge to date has been to identify if and where pings are being blocked or rate limited. Blocking or rate limiting is increasing, especially for developing countries. Blocking is relatively easy to detect while evidence of rate limiting is more subtle. The onset of rate limiting is usually observed as a step increase in the losses for a site, while neighboring sites may see no such increase. A test for ping rate limiting can be performed by opening and closing TCP connections [4] to a well known port (e.g. web port 80) on the Remote-host. The Monitoring-host sends a SYN packet to open a TCP connection. The Remote-host responds with an ACK packet acknowledging the connection. The Monitoring-host then closes the connection. If many fewer ACKs are lost than ping echo response packets, then it can be deduced that the pings are being rate-limited. In such cases, a `traceping` (a `traceroute` followed by the pinging of all the routers along the route) may be useful in ascertaining where the rate limiting starts. In cases where the limiting occurs close to the remote site, it may be possible to contact the site and request pings to be allowed between the Monitoring- and Remote-hosts. When the blocking is not at the site but at an intermediate ISP (for example the Australian AARnet and the Vietnamese

VISTA), and there is little success in contacting the intermediate ISP, the Remote-host is removed from the monitoring. This was the case when the SYN/ACK tests were used to check sites in Australia. It turned out that all the attempted Australian sites were being rate limited in Seattle.

Continuing Efforts

While funding exists, continual efforts are being made to maintain, upgrade, and expand the PingER deployment. With respect to the Remote-hosts, the primary effort is to keep them accessible to the Monitoring-hosts. For Monitoring-hosts, it is important to keep the lists of Remote-hosts up to date, the monitoring code running reliably and the data accessible for collection. For the Archive/Analysis-hosts, the continuing effort is to assure that the data is reliably collected and archived from all the Monitoring-hosts and analyzed. The data, analysis results and reports are made available to all interested users via the data selection, analysis, and display tools.

Conclusion

PingER provides a valuable light-weight tool for the active end-to-end performance monitoring of networks around the world. With its continual gathering of internet monitoring data, PingER provides extensive quantitative historical and near-realtime analytical information on world-wide networks. PingER has shown itself to be useful for recognizing developing network problems, troubleshooting network problems, gauging repetitive network behaviors, choosing ISP's and most importantly, providing valuable quantitative information useful for qualifying network needs and improvements.

Acknowledgements

References

[1] W. Matthews & R. L. Cottrell, "The PingER Project: Active Internet Performance Monitoring for the HENP Community", IEEE Communications Magazine, May 2000.
 [2] "Virtual Silk Highway", available at http://www.silkproject.org/
[3] R. L. Cottrell, "Measuring the Digital Divide with PingER", available at http://www.slac.stanford.edu/grp/scs/net/talk03/ictp-oct03.ppt2
[4] "Synack", available at http://www-iepm.slac.stanford.edu/tools/synack/