

What is the Internet Doing for and to You?*

R. L. A. Cottrell

Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309

C. A. Logg

Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309

D. E. Martin

HEP Network Resource Center, Fermi National Accelerator Center, Batavia, IL 60510,
USA

Many HEP sites have been participating in a project to monitor end-to-end Internet performance. A few sites have acted as central collection, analysis and report generation points for the large amount of data generated. The reports are useful for short term trouble identification, long term projections, and the evaluation of Internet services. The work is illustrated by examples of both detailed and more general Internet performance.

Presented at Computing in High Energy Physics CHEP97, Berlin Germany, April 7-11, 1997.

*Work supported by Department of Energy contract DE-AC03-76SF00515.

What is the Internet Doing for and to You?

R. L. A. Cottrell ^a, C. A. Logg ^a, D. E. Martin ^b

^a *Stanford Linear Accelerator Center, Stanford, CA 94309, USA*

^b *HEP Network Resource Center, Fermi National Accelerator Center, Batavia, IL 60510, USA*

Abstract

Many HEP sites have been participating in a project to monitor end-to-end Internet performance. A few sites have acted as central collection, analysis and report generation points for the large amount of data generated. The reports are useful for short term trouble identification, long term projections, and the evaluation of Internet services. The work is illustrated by examples of both detailed and more general Internet performance.

Keywords: Wide Area Networking; Network; Monitoring; End-to-end; Performance

1 Introduction

Many concerns are being raised in research and education communities about poor Internet performance impeding the ability to participate effectively in distributed collaborations. This is particularly significant in the HEP community with large collaborations of hundreds of scientists in many countries working on major experiments. To understand these concerns the Energy Sciences network (ESnet) Site Coordinators Committee (ESCC) chartered a task force to investigate tools and techniques for monitoring traffic, reliability and consistency between ESnet sites and sites of interest in the Internet. In 1996, the group coalesced around work begun at the Stanford Linear Accelerator Center (SLAC) [1].

2 End-to-end Measurements

These measurements focus on using the Internet ping utility. The rationale for using ping and the measurement methodology are discussed in a companion paper [2].

RAL (England) – Jan. 30th, 1997




Fig. 1. Average, minimum and maximum (of ten 1000 byte payload pings) response time between SLAC and the Rutherford Appleton Laboratory (RAL) in England, for one day. The time of day is in Pacific time, so midnight on the graph is 8 am in England.

The data is analyzed and presented in reports accessible via the WWW. The reports provide both short term (last few hours or days) and longer term (last fortnight, last 10 weeks, last 180 days and going back several years) information. The short term reports are mainly used for trouble-shooting and understanding the current state of connectivity. The longer term reports are mainly for looking at trends and planning.

2.1 Short Term Reports

Short term reports provide information on the performance (response time, packet loss, reachability) of remote hosts measured so far this day, for yesterday, for the last 14 days, and for the last 30 days. Figure 1 shows an example of such a report for one day. The day versus night (at RAL) response time differences are striking in this figure and indicative of congested links during the English daytime.

To generate alerts we calculate the averages and standard deviations of the ping response and packet loss to each remote host for the previous 10 weeks. This is repeated for the last 7 days and for yesterday. Alerts are raised if the last 7 days or yesterday's averages are over 3 standard deviations greater than for the last 10 weeks. These alerts may be provided via email or active links in WWW reports.

2.2 Medium Term Reports

Medium term trends are provided by plots showing the average daily ping response and packet loss for the last 180 days for each remote host. Figure

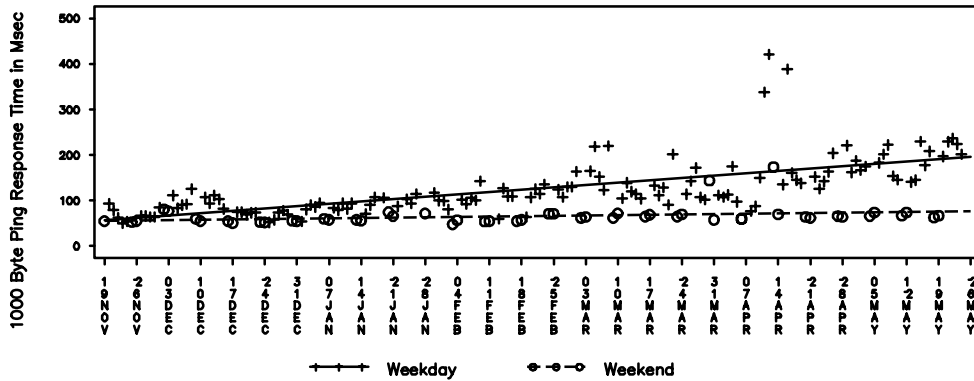


Fig. 2. 180-day trend plot of the ping response (for 1000 byte ping payload) between SLAC and UCD starting 19 November 1995. The lines are linear regression fits to aid the eye.

2 shows a typical example measured between SLAC and the University of California at Davis (UCD). Immediately visible in this figure is a degradation in weekday response by almost a factor of 4 in this 180 day period. The packet losses (not shown here) for this period are more variable and increase by about 50%. The big difference between weekday and weekend performance is again indicative of congestion.

To provide a measure of performance predictability (in particular the variability between day and night time performance), we calculate for each set of 30 minute interval measurements the 100-byte payload packets the ping $success = 1 - \text{packets lost}^1 / (\text{total number of packets})$, and for 1000-byte payload packets the ping $data_rate = 2000 \text{ bytes} / (\text{average response time of 10 consecutive ping packets})$. Then for all $successes$ and $data_rates$ to a given host in one day we calculate the dimensionless ratios: $s = avg(success) / max(success)$ and $r = avg(data_rate) / max(data_rate)$. We then scatter plot the daily success ratio (s) versus the rate ratio (r) for a given month for a set of hosts. Values of the ratios close to one indicate the average performance is close to the optimal performance. Ratios much less than one occur particularly on links which are congested during prime time. Examples of such plots are seen in Figure 3. which indicate that performance predictability was much better between SLAC and ESnet hosts than between SLAC and western North American hosts, presumably because SLAC has a better connection to ESnet hosts without passing through the commercial Internet.

¹This excludes measurements with 100% packet loss, these are accounted for in the *unreachability* analysis.

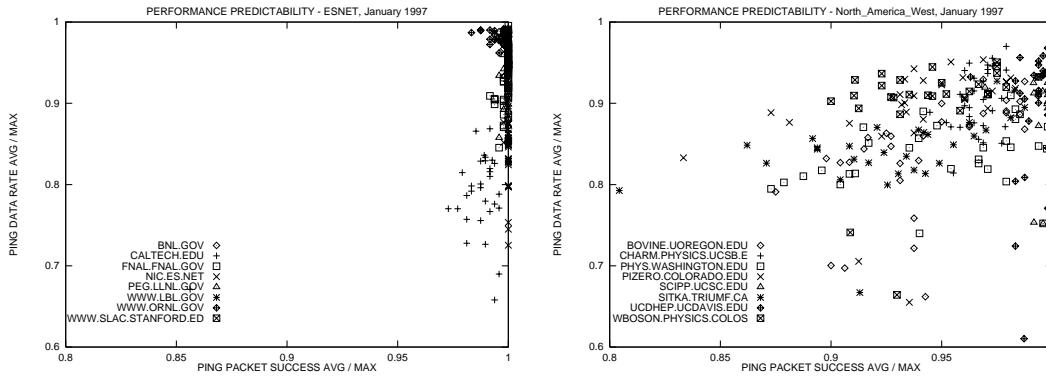


Fig. 3. Performance predictability between SLAC and ESnet hosts and between SLAC and western N. American hosts monitored by SLAC for January 1997. Each point on the plots represents (s, r) for one host for one day.




Fig. 4. Average SLAC prime time (7am - 7pm, weekday) monthly packet loss between SLAC and some western N. American HEP hosts. Note the general improvement Apr-Jun '96 following the improvement in ESnet connections to the Internet. Also it is seen that there is considerable variability from month to month and host to host.

2.3 Long Term Reports

To provide trends going back over longer periods, we calculate the average response time and the average ping loss for each month for each host. Since most of the interest concerns the performance during working hours, we include only weekday ping data measured between 7am and 7pm (SLAC time). An example of such a plot is seen in Figure 4.

We identify a host as being down or unreachable when no ping response is obtained in the set of 21 pings made each 30 minute period. Using this identification, we calculate the ping *unreachability* = $(\# \text{ periods host is down}) / (\text{total number of periods})$, the number of down periods per month, the Mean

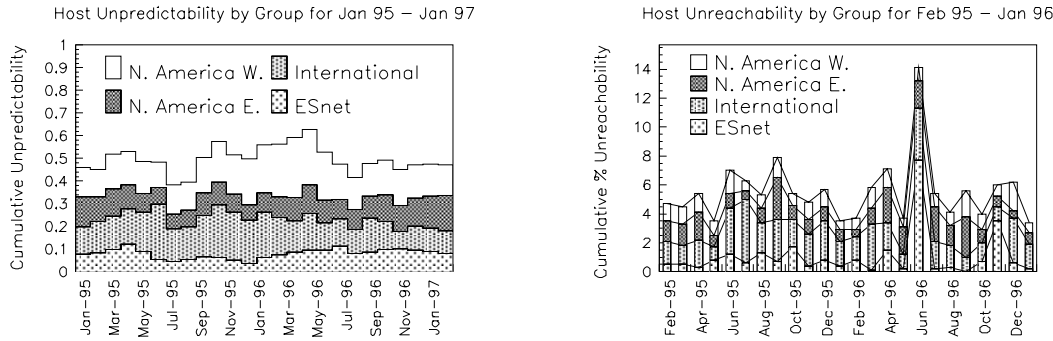


Fig. 5. Group ping unpredictability and unreachability. It can be seen that the unpredictability is worse for international hosts and western N. American hosts. Western N. American hosts became particularly unpredictable in Spring 1996 before the ESnet links to the Internet were improved. The unreachability peak in June 1996 was due mainly to a host which was shutdown while it was relocated.

Time Between Failures for each remote host. The unreachability is plotted (one point per host/month) and the other information is provided in tabular form.

To provide a broader overview of the performance, we average the various indicators (response, packet loss, unreachability, unpredictability (defined as the distance of the coordinate (s, r) from $(1,1)$)) for each month over groups of hosts. Typically the grouping is geographical or by service provider. The main host groupings we use are: ESnet, western N. America, eastern N. America, international (non ESnet and non N. American), and local Internet Service Providers (ISPs) in the San Francisco Bay Area. Figure 5 shows examples of group ping unreachability and unpredictability.

One of the more easily understood and critical metrics, for the end user, is the packet loss, since packet loss results in timeouts which have a large impact on the performance of network applications. To provide an upper level view of the packet loss we arbitrarily divide the losses into 5 *quality* categories:

- $\leq 1\%$ ² packet loss \equiv Good WAN performance.
- > 1 & $\leq 5\%$ packet loss \equiv Acceptable WAN performance.
- $> 5\%$ & $\leq 12\%$ packet loss \equiv Poor WAN performance.
- $> 12\%$ ³ & $\leq 25\%$ packet loss \equiv Bad WAN performance.
- $> 25\%$ packet loss \equiv Unusable WAN Performance.

We then find the percentage of months for which each host fell in each of the

² 1% is the threshold we use on the SLAC local area network for generating alerts.

³ The "Internet Weather Report" (<http://www.internetweather.com/>) marks networks as "RED" if the packet loss is $> 12\%$.




Fig. 6. SLAC prime time ping loss quality distributions for groups of hosts from January 1995 through December 1996. The numbers in parentheses are the number of host-months and the median packet loss.

above categories. We average these percents for each group of hosts and plot the distributions (i.e. host group vs. category vs. percentile in each category) which is shown in Figure 6. ESnet ping loss performance is seen to be good or acceptable over 95% of the host-months. The other groups, however, typically have packet losses which are poor or worse over half the host-months.

3 Conclusions and Futures

Ping is an excellent tool for end-to-end network performance monitoring. It provides almost universal coverage. Administrators at the monitored remote hosts do not have to install any software. It has low network impact if used wisely. It provides useful short- and long-term measures of bottleneck bandwidth, available bandwidth [3], response time, packet loss, reachability, and predictability which can be related to user applications. A major challenge has been coming up with simple, intelligible ways to characterize and visualize the enormous amounts of data.

The results indicate that by most measures, performance within ESnet is acceptable to good. However packet loss performance between ESnet and the Internet at large is, on average, poor or worse for the hosts monitored. Packet loss seen from SLAC for non-ESnet hosts improved dramatically between April and June 1996, and the improvement has been sustained. In general performance is very variable in both the short and long-term, particularly for international hosts. From SLAC, average monthly response times by host groups are typically 300-500 ms. for international hosts, 150-220 ms. for eastern N. American hosts, 80-140 ms. for western N. American hosts, and 40-50 ms. for ESnet hosts.

The methodology is also being utilized to: select ISPs and monitor their performance possibly with a view to writing a service contract; help decide which

universities to connect directly to ESnet; and, to identify bottlenecks in order to decide where to focus efforts.

Possible future work includes: performing the measurements with Poisson sampling, which, in principle results in unbiased measurements, even if the sample rate varies [4]; looking at better definitions of prime hours so the definition is less monitoring site specific and more realistic for international links; increasing the number of monitoring sites by distributing the monitoring software (in particular add monitoring hosts in other countries); responding to requests by a number of HEP-related organizations to add more remotely monitored hosts; looking at how to combine the data and analyses done at the analysis sites; consider a more careful choice of remote hosts to monitor (more lightly or uniformly loaded, and pings are responded to at a high priority - e.g. in the kernel); and, installing a range of fixed size WWW pages at various sites to look at long term WWW responsiveness.

4 Acknowledgements

We would like to thank Bill Wing for help and encouragement in this work, Lois White and Charley Granieri for help with producing graphs, John Halperin for expert criticism, and Mike Wendling for much help with reformatting plots for embedding in LaTeX.

References

- [1] C. A. Logg and R. L. A. Cottrell *Network Monitoring and Performance Management at SLAC*. SLAC-PUB-95-6744, Stanford Linear Accelerator Center, CA. Presented at Networld+Interop Engineers' Conference, Las Vegas, March 1995.
- [2] D. E. Martin, R. L. A. Cottrell, and C. A. Logg *Internet Performance and Reliability Measurements*. Talk at CHEP97.
- [3] R. L. Carter and M. E. Crovella *Measuring Bottleneck Link Speed in Packet-Switched Networks*. Tech. Report BU-CS-96-006, Computer Science Department, Boston University, March 1996.
- [4] V. Paxson *End-to-End Routing Behavior in the Internet*. Proc. SIGCOMM '96, August 1996.