

## Just How Good are You? Measuring Network Performance

February 2003

**Geoff Huston**

If you are involved in the operation of an IP network, a question you may hear often these days is: "How "good" is your network?" Or, to put it another way, how can you measure and monitor the quality of the service that you are offering to your customers? And how can your customers monitor the quality of the service you provide to them? How can you describe and measure the performance of an Internet network?

Of course if you are a customer of an Internet Service Provider the same question holds. How "good" is the service, and how can providers' service performance be compared on an equal basis?

These questions have been lurking behind many public and enterprise IP networks for many years now. With the increasing levels of deployment of various forms of high speed broadband services within today's Internet there is new impetus to find some useable answers that allow both providers and customers to place some objective benchmarks against the service offerings. With the lift in access speed with broadband services there is an associated expectation on the part of the end user or service customer about the performance of the Internet service. A higher speed service should be "better" in some fashion, where "better" relates to the performance of the network and the service profile that is offered to network applications. And not only is there an expectation of "better" performance in somewhat nebulous terms, it should be a measurable artifact of the service.

As well as technical motivations, there is also a business driver behind this question of how to measure a network's performance. In increasingly significant segment of the Internet market is moving away from a simple cost-driven undifferentiated commodity market into a market where there is a serious attempt to provide differentiation based on the performance of the delivered service. So it seems that performance measurement is now an important aspect of any ISP's operation. But what is Internet performance? How can you measure it?

An informal functional approach to a definition of network performance is measuring the speed of the network. How fast is the network? Or, what's the elapsed time for a particular network transaction? Or, how quickly can I download a data file or a web page? This measurement of time for a network transaction to complete certainly relates to the speed of the network, and speed is a good network performance benchmark, but is speed all you need to measure?

When looking at the broad spectrum of network performance, the answer is that speed is not everything. The ability of a network to support transactions that include the transfer of large volumes of data, as well as supporting a large number of simultaneous transactions are also part of the overall picture of network load and hence of network performance. Here a network must provide adequate bandwidth to provide throughput for these transactions. The IP protocol suite uses the Transmission Control Protocol (TCP) for managing such data transfer, and with TCP there are a number of network attributes that impact on TCP performance in addition to network bandwidth.

TCP uses a control feedback loop between the sender and the receiver, and, as a general rule, the lower the time lag for the feedback system the more accurately TCP can adapt to the constantly changing network conditions and operate efficiently. Therefore end-to-end delay, or latency, is an important consideration for network performance.

But handling large data sets is not everything in performance. Consideration should also be given to the class of network applications where the data is implicitly clocked according to some external clock source. Such real time applications include interactive voice and video, and their performance requirements include the total delay between the end points, or latency, as well as the small scale variation of this latency, or jitter. Performance measurements also include the ratio of discarded packets to the total number of packets sent, or loss rate, as well as the extent to which a sequence of packets is reordered within the network, or even duplicated by the network. Taken together, this set of performance factors can be considered as a form of the amount of distortion of the original real time signal.

Accordingly, a functional description of network performance encompasses a description of speed, capacity, latency and distortion of transactions that are carried across the network. This informal description of what constitutes network performance certainly feels to be on the correct path, given that if one knew the latency, available bandwidth, loss and jitter profile and packet reorder probability as a profile of network performance between two network end points, as well as the characteristics of the network transaction, it is possible to make a reasonable prediction relating to the performance of the transaction.

Of course the tricky part is working out how to measure these quantities and then map them back to an overall picture of network capability and performance.

Its here that service providers and customers often find themselves with entirely different motivations in service performance measurement. The service provider wants to measure the quality of the network itself. Normally this would relate the measurement of a transit path, commencing when a packet enters the provider's network, and taking the measurement outcome as the packet leaves the provider's network. The customer, on the other hand has less of an interest in the performance of the network, and more of an interest in the performance of the application itself, spanning the entire path from the client to the server and back again. IN the context of the Internet, such paths may transit a number of provider's networks, and its the cumulative picture rather than the profile of any individual network that is of interest to the customer. The service provider also measures different aspects of performance. As we've noted already, the service provider is interested in the per packet transit latency and the stability of the latency readings, the packet drop probability and the jitter profile. The end user has a somewhat different, and perhaps more fundamental set of interests: Will this voice over IP call have acceptable quality? How long should this download take?

So what tools exist to help us to measure network performance?

Many network performance management systems and customer performance management systems are based on a very simple tool: *ping*. The measurement system sender generates an IP ICMP echo request packet, and addresses it to a target system. As the packet is sent, the sender starts a timer. The target system simply reverses the IMCP headers and sends the packet back to the sender as an ICMP echo reply. When the packet arrives at the original sender's system the timer is halted and the elapsed time is reported.

Ping is simple, efficient, widely used, and for network performance measurement, often terribly misleading. In measuring the elapsed time from the application sending the packet to the application receiving a matching response, there are a number of variables, including the

granularity of the sending system's clock, the scheduling algorithm used by the sender and the relative priority of the measurement application, the load on the target system and the relative scheduling priority given to responding to ICMP requests, all added to the transit time to send the packet through the network and the time to send the matching response. Surely we are talking only milliseconds? True, but in high speed networks where a transcontinental delay is only tens of milliseconds and jitter is sub-millisecond, then these additional sources of delay become a real factor in masking the true network measurement.

As a performance diagnostic tool, ping is a relatively coarse and insensitive instrument. Can we do better? Yes, certainly. One of the most promising approaches in the One-Way synchronized measurement.

The One-Way approach does not use a single network management system, and a set of targets, but relies on the deployment of a collection of probe senders and receivers using synchronized clocks. This moves beyond a simple and ubiquitous software tool that everyone, providers and customers alike can run, into a specialized environment that is specifically configured to measure the characteristics of particular network transit paths with very high accuracy.

The One-Way methodology is relatively straightforward. The sender records the precise time a certain bit of the probe packet is transmitted into the network; the receiver records the precise time that same bit arrives at the receiver. The two clocks have to be in sync, and achieving this to microsecond accuracy is an interesting problem. Initial implementations of this approach have used Global Positioning System satellite receivers as a synchronized clock source. One of the noted problems with the use of GPS was that computers are generally located within machine rooms and a clear GPS signal is normally only available on a rooftop. Later implementations of this approach have used the clock associated with the CDMA mobile telephone network as a highly accurate synchronized distributed clock source, with the advantage that the time signal is usually available close to the measurement unit. Consequent correlation of the sender's and receiver's data from repeated probes can reveal the one-way delay and loss patterns between sender and receiver.

For the service provider this system can provide a very accurate view of the behaviour of the active network elements within a select set of network transit paths, using the metrics of latency, jitter and loss as described above. As a real-time diagnostic tool it can allow a network operator to maintain a constant view of network behaviour and complement active polling as a means of managing the network.

But can this help the customer in assessing the performance of their provider? There is some potential here, depending on how the reporting relationship is phrased between the provider and the customer. While many forms of performance reporting involve reporting on various averages of latency and loss, there is also the capability of providing more detailed data as a real time feed.

There are other ways of manipulating ping to provide more information. One way is to vary the size of the packet. Larger packets take a longer amount of time to be passed along a constant size transmission path, and by comparing the latency times of various sized packets it is possible to build up a picture of the capacity of a transmission path using ping as a remote probe. Another method of ping manipulation is to vary the sending rate of the ICMP packets. If a TCP flow control algorithm is used to control the sending rate of ping packets it is possible to infer the likely TCP peak data transfer rate between two points.

But perhaps in this we have lost sight of the original objective here, and it may be useful to return to the original question of how we can measure the performance of an IP network. The above techniques allow individual paths within a network to be measured for various characteristics, and, with some approximation, these results can relate to some measure of network performance. But

there is still a sense that there is something missing in all this. How "good" a network is, from a user's perspective, is equivalent to how well applications perform across the network. The basic Internet architecture is one of end-to-end data flows, where the network's task is one of simple packet switching. The Internet architecture does not manage the network resource by trying to 'protect' one application's use of the network from any other. This task is left to the application to attempt to sense the current state of use of the network and adapt its own demands to that of a fair share with all other applications who are undertaking a similar adaptation of their own. How can a network provider measure this form of adaptive cooperative behaviour and create a performance metric? Unfortunately this is a somewhat challenging question, and one without clear answers so far. One thing we do know, is that this is not an issue unique to the Internet. Measuring the performance of a metropolitan road network has similar properties of attempting to relate adaptive traffic components along specific paths to the performance of the system as a whole.

## Further Reading

There is a wealth of material on the Internet on the topic of network measurement, and the major exercise is undertaking some filtering to get a broad collection of material that encompasses a range of perspectives on this topic. These sources have been used to prepare this article, and are recommended as starting points for further exploration of this topic.

Internet Performance Survival Guide , Geoff Huston, Wiley Computer Publishing, 2000.

IPPM Metrics for Measuring Connectivity, J. Mahdavi, V. Paxson, [RFC 2678](#), September 1999.

A One-way Delay Metric for IPPM, G. Almes, S. Kalidinki, M. Zeukuaskas, [RFC2679](#), September 1999.

A One-way Packet Loss Metric for IPPM, . Almes, S. Kalidinki, M. Zeukuaskas, [RFC2680](#), September 1999.

The RIPE Test Traffic Measurement service at <http://www.ripe.net/ripenc/mem-services/ttm/>

Treno , online at [http://www.psc.edu/networking/treno\\_info.html](http://www.psc.edu/networking/treno_info.html)

Trends in Measurement and Monitoring of Internet Backbones, session at the 26th North American Network Operators Group, hosted by D. Meyer, <http://www.nanog.org/mtg-0210/measurement.html>, October 2002

Some thoughts on CoS and Backbone Networks, D. Meyer, presentation to the IEPREP Working Group, IETF-55, <http://www.maoz.com/~dmm/IETF55/ieprep/>, November 2002

Netflow resource page:  
[http://www.cisco.com/warp/public/732/Tech/nmp/netflow/netflow\\_techdoc.shtml](http://www.cisco.com/warp/public/732/Tech/nmp/netflow/netflow_techdoc.shtml)

Netramet, and many other interesting measurement tools are referenced in a resource page at <http://www.caida.org/tools>

This is also an active area of research and there are a number of activities that are in the area of research group activities and workshops.

The Internet Research Task Force has an Internet Measurement Research Group. Further details can be found at <http://www.irtf.org/charters/imrg.html>

ACM SIGCOMM, the ACM Special Interest Group on Data Communications sponsors an Internet Measurement Workshop. Proceeding of the November 2002 workshop can be found at <http://www.acm.org/sigcomm/imw2002/>

The details of the 2003 Passive and Active Measurement Workshop can be found at <http://www.pam2003.org>.

---

## Disclaimer

The above views do not represent the views of the Internet Society, nor do they represent the views of the author's employer, the Telstra Corporation. They were possibly the opinions of the author at the time of writing this article, but things always change, including the author's opinions!

---

## About the Author

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for the past decade, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. Huston is currently the Chief Scientist in the Internet area for Telstra. He is also a member of the Internet Architecture Board, and is the Secretary of the APNIC Executive Committee. He was an inaugural Trustee of the Internet Society, and served as Secretary of the Board of Trustees from 1993 until 2001, with a term of service as chair of the Board of Trustees in 1999 – 2000. He is author of *The ISP Survival Guide*, ISBN 0-471-31499-4, *Internet Performance Survival Guide: QoS Strategies for Multiservice Networks*, ISBN 0471-378089, and coauthor of *Quality of Service: Delivering QoS on the Internet and in Corporate Networks*, ISBN 0-471-24358-2, a collaboration with Paul Ferguson. All three books are published by John Wiley & Sons.

E-mail: [gih@telstra.net](mailto:gih@telstra.net)