



Making Sense of Broadband Performance

**Never Mind the Connection Speed,
Measure the Connection Quality!**

*In today's Internet world dominated by on-demand Voice,
Video and IP-TV, the key to good connection performance is
consistent throughput – not fast speeds.*

Never mind the connection speed, measure the quality!

Knowing your Internet speed is slow without providing the means to help resolve the problem is of little value.

When discussing Internet connectivity problems the number one complaint is always about speed or the lack of it. End-users relay the same message over and over again, and quite frequently users with Internet service problems are quick to defend their choice of service provider with statements like “But I have a 5 Mbps or a 10 Mbps connection so I shouldn’t have jerky video” or “But the speed tester I use says I can run 9 VoIP lines, so why are my VoIP calls often garbled?”

As *speed* is the prime criteria when users select an Internet service provider, when they are asked “what is the quality of the service?”, the question is most often met with a puzzled along with the retort, “Quality! What do you mean by quality?”

Speed versus Quality is a crucial issue for measuring bandwidth performance. An example of speed without quality would be like driving a Ferrari during the heavily congested rush hour – while the capability of high speed exists, the speed is not realized due to other factors.

We all drive on highways and through experience we know only too well that our journey time will not be governed by the maximum speed of our car or the maximum speed of the highway. In reality it will be governed by the many different events that occur on our journey such as weather, congestion, accident backup or highway construction.

Running a speed test that shows a slow throughput without consideration of the event information to help the ISP to identify and resolve inherent performance issues has absolutely no value. Only through careful measurement of every aspect of the end-to-end journey can a speed test really start to benefit the user, and more importantly benefit the service provider by providing the information needed to identify the problem.

In the Internet world there are many different issues that affect how our applications perform. When performance is as we expect we are content, when it is not we get frustrated. Understanding the *events* that impact your Internet connection along with the *quality* of service delivered for your applications is vitally important for the problems to be resolved. In our highway example, if the quality of the highway service was to be measured by its users, it would be a measure of the highway’s ability to deliver you to your destination in a time that closely matches the time it would take if you were able to sustain the

maximum allowed speed. In other words, assuming you are driving as fast as allowed by the legal limit, the calculation of highway quality would be:

$$\text{Highway Quality \%} = \frac{\text{Your Speed}}{\text{maximum allowed highway speed}} \times 100$$

If the highway is legally limited to 65 mph and your average speed is 35 mph then the quality is:

$$\text{Highway Quality \%} = \frac{35}{65} \times 100 = 53\%$$

The quality measure is critical because if the highway quality is such that you are only able to maintain 30 mph when the application (arriving at the airport to catch the flight) requires 35 mph, you will certainly miss your flight.

I know how to measure speed, how do I measure the quality?

The key to success of an Internet connection is a combination of a good speed with a good quality of service. In fact, it is preferable to have a slower 3 Mbps (Megabits per second) connection with a 99% quality of service rather than a 6 Mbps connection with a 50% quality of service. Both will achieve about the same throughput overall, however the delays inherent in the packet flow that result in the lower of quality of service will adversely impact time-dependant applications such as VoIP, video or MP3.

In order to measure quality, an Internet connection testing application has to measure the *events* for the Internet connection as well as the speed. Otherwise, you only know that the traffic is slow but you have no clue as to the cause. To better ascertain the cause requires that the testing application is capable of measuring the *traffic flow* as well as the *traffic speed*. Graphical views of the traffic flow show the inherent delays occurring with the data along with the all important timing.

Once the *delay* picture is exposed, you can start to see the impact the delays have on the data movement and thusly on the applications. For example, in line with our analogy of driving to the airport, you will be able to see those parts of the data packet journey that are fast and those parts that are not. You will also be able to see if the fast data movement periods are reaching the contracted maximum or whether, as is often the case, the fast periods are still not meeting the contracted service levels. In addition, the volume of delays, size of the delays and the timing of the delays provide the insight to understand whether the problem is *traffic* related (i.e. Internet congestion), or whether it is *police* related, i.e. some other sort of data traffic management or time multiplexing is being imposed by the ISP. The chart below shows a good example of speed versus quality.

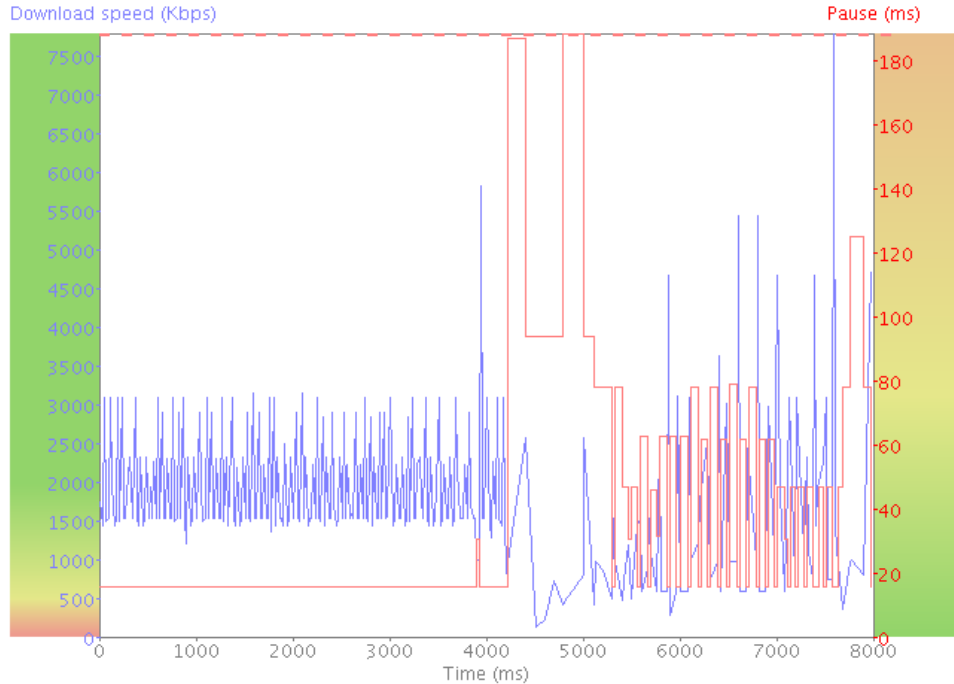


Fig. 1 – MySpeed® chart showing data transfer delays

In the chart plotted above the TCP delay (red line) suddenly impacts the data speed test at about 4 seconds into the test. The throughput or bandwidth speed (blue line) drops from 3 Mbps to approximately 0.2 Mbps before recovering and compensating some 2 seconds later with throughput at 3 Mbps and spikes to 5, 6 and even 7 Mbps. A conventional speed tester would most probably have indicated an average speed of around 3.5 Mbps. The measurement would most probably be accepted as *OK* by the tester as the spikes elevate the test to be close enough to the contracted service of 4 Mbps. In reality, there are a number of issues that need to be addressed because the sudden drop was severe and also sustained (2 seconds). Note that with peak bandwidth showing at 5 Mbps to 7 Mbps the connection is able to compensate for the drop with higher data speeds to bring the average up to 3.5 Mbps or more. Unfortunately the average speed is fine but the quality is poor. It is likely that the quality of a VoIP call or video session would suffer under such erratic and unpredictable conditions with TCP delays exceeding 180 ms (milliseconds) or more.




Fig. 2 - MySpeed chart showing data transfer delays

The longer the time delay between data transfers, the more dramatic the impact on service quality. In Fig. 2 it would have been very difficult to maintain any quality of service as the delay was almost a full 3 seconds and the data throughput dropped to below modem speeds for the duration.

What should a good quality connection look like?




Fig. 3 - MySpeed chart showing good connection quality

Interestingly enough, the two charts (Figures 2 and 3) above were taken from the same office location but from two different PCs connected to different ISPs. The distance to the testing server was identical, just the ISPs were different. The first chart was a 4 Mbps connection and the second was a 1.5 Mbps connection. In the case of the latter (Fig. 3) you can see the service is very clean and reflects a constant and consistent data flow. The delays are a consistently low 7 ms and even the peak delays are only 15 ms.

Understanding the nature of the delays such as timing, size and frequency provides a good clue as to the cause. The two prime causes are *traffic congestion* and *traffic control*. Congestion is a very common problem because ISPs often over subscribe circuits in the local exchanges. Control is also common as ISPs divide high bandwidth connections using time-sharing algorithms. In other words, you think you have a 5 Mbps connection when in reality you have a regulated 5 Mbps slice of a 100 Mbps connection. This is a common approach of cable providers.

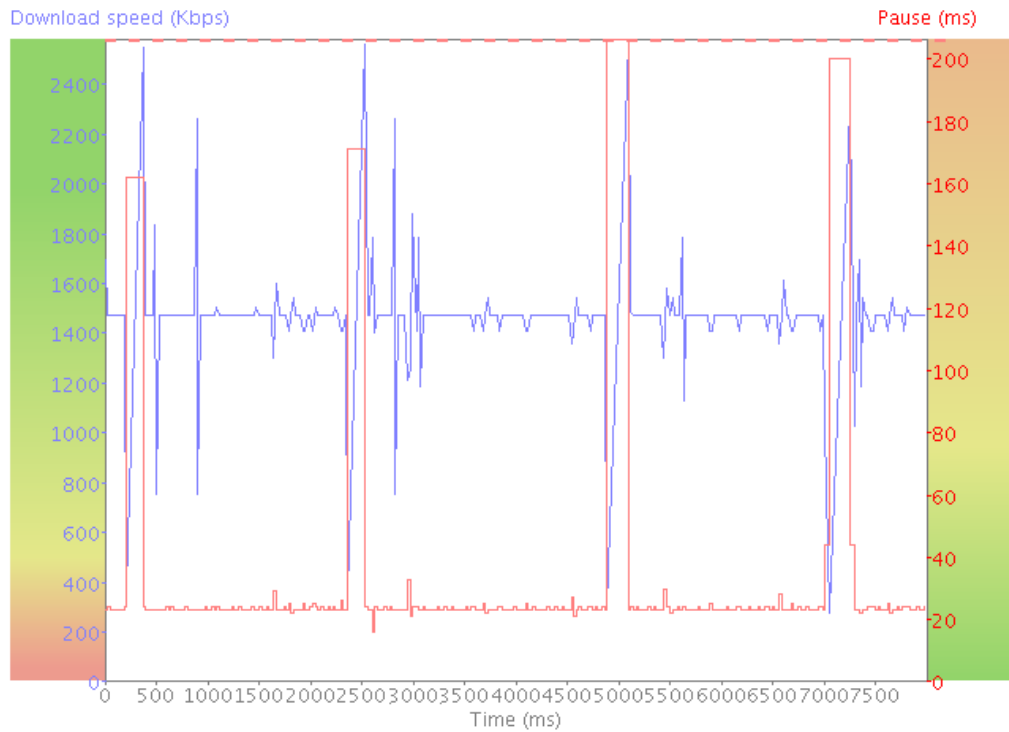


Fig. 4- MySpeed chart showing *regulated* throughput

If the delays are man-made, or by design rather than random chance, then a regular pattern will quickly show. In Fig. 4 above there is a 200 ms delay every 2 seconds. You can see the impact of that delay on the data flow and with it the impact on applications, especially time sensitive applications such as VoIP or video.

In terms of throttle-based bandwidth management as shown above, to a large extent the impact will depend on the intelligence and sophistication of the implementation. For example, many small delays versus fewer large delays will deliver different characteristics in the data flow.




Fig. 5- MySpeed chart showing poor *regulated* throughput

Fig. 5 above shows a more dramatic impact to the data flow over Fig. 4 as speeds degrade to below 200k. Notice that in most cases the data flow also records a higher than normal spike after the delay, which improves the average speed performance, but not necessarily the performance quality. This often leads the user to believe that the speed is acceptable, but then cannot understand why VoIP calls keep dropping.

In extreme cases staggering TCP delays are being recorded, and the user still believes that the speed average is acceptable.




Fig. 6 – MySpeed chart showing dramatic delay impact

Fig. 6 above shows such an extreme case. During the first half if the test (about 4 seconds) the average throughput recorded just exceeded the contracted 1.5 Mbps. In the last half of the test it was less than 0.5 Mbps. Is that a good 1.5 Mbps throughput? In reality the connection shows 10 Mbps peaks with 0.0 Mbps (zero) troughs. This is not a quality connection by any account and is incapable of sustaining time-dependent applications.

In short, knowing the speed of a connection is only a small part of the picture, and does nothing to identify a connection performance problem. The larger part of the picture is the consistency of the data flow – connections with wide variations in speeds or large gaps between data transfers will cause applications like VoIP and video to perform poorly. To truly understand connection performance, good or bad, it is essential to measure the quality.

Web applications need a quality connection - never mind the speed measure the quality!

For information on how MySpeed can help you measure connection quality, please see www.myspeed.com.

Visualware Inc.
PO Box 668
Turlock, CA 95381
866-847-9273
www.visualware.com