

Should I trust my PingPlotter results? (Definitely)

Saved From: <https://www.pingman.com/kb/article/should-i-trust-my-pingplotter-results-definitely-46.html>

Question

I sent my PingPlotter graphs and information to our network administrators and they responded that traceroute is not an accurate way of troubleshooting networks. They say that the first slow hop is the weakest link and invalidates any other data reported. They also say that ICMP traffic (like PingPlotter creates) is low-priority and can't be trusted. They say they can't trust PingPlotter data, so stop sending it to them.

Solution

These statements have enough truth in them to cause a lot of users to leave a network administrator alone, but are targeted to do this - drive an end-user to stop bringing these results.

Both latency and packet loss (the two things that PingPlotter excels in capturing, measuring and displaying) are additive. This means that a slow connection early in your route will, indeed, add latency at all downstream hops. The part left off the 'objections' statement you got is that any additional latency or packet loss problems can still be seen - even if your first hop is a modem. If hop 1 adds 150ms and 5% packet loss, you can still see that some other hop (say, hop 7) adds another 100ms of latency and another 10% of packet loss. The fact that there is some latency and packet loss at hop 1 does not make the information gathered about hop 7 any less valid. PingPlotter is especially useful here because you can sample enough times to make the differences in packet loss and latency at different hops statistically valid. The difference between 2% and 4% packet loss needs at least 500 samples to be statistically convincing - the stock traceroute utility does 3 samples, while PingPlotter allows you to do thousands or more.

In reality, you're shooting for 0% packet loss and a reasonable latency across your entire route. When you see packet loss in your route, you almost certainly want to resolve that problem. If the packet loss appears at a link you know is slow, it still needs to be explained and possibly resolved. If it appears at some other hop, that needs to be investigated too - just because one of your early hops is slow doesn't make that packet loss reading any less important to solve.

The 'weakest link' isn't the right way to see things. A weakest link means that it will be the first thing that breaks and you won't be able to see any information about any other links. With PingPlotter, you certainly see the latency and packet loss characteristics for a weak link, but if there's another link (a bit less weak), you also see the performance characteristics for that link, and any other link.

As for ICMP accuracy, again this is not entirely wrong. There are certainly routers out there that down-prioritize ICMP traffic under heavy load. However, most of the routers on the internet route ICMP just fine - especially when not under heavy load. Most routers do prioritize ICMP traffic below HTTP (or, voice traffic, or other kinds), but a well-running network is running below capacity. At 50% capacity, these prioritization decisions don't have to be made in a way that significantly impacts your measurements. If routers do need to start making decisions like this, then it's because they are reaching capacity, and there's a fine line between dropping/delaying ICMP traffic and dropping other kinds of traffic - so in many cases, when a router decides it needs to start dropping ICMP traffic, it's also delaying other kinds of traffic. This is exactly the scenario we want PingPlotter to capture! The fact that routers down prioritize ICMP means that PingPlotter traces are often an early indicator of problems.

PingPlotter also offers additional packet types - you can use both TCP and UDP packets for networks that block ICMP. This gives you the capability of seeing the final destination's latency and packet loss using the same (or very similar) packet types that your services use. Intermediate hops still rely on ICMP TTL expired packets, but the additional packet types let you see statistics about the final destination and then use [PingPlotter troubleshooting techniques](#) to evaluate if the intermediate hops are reporting valid data.

You do always need to correlate PingPlotter data with some other network results. The fact that ICMP traffic doesn't always 100% match other kinds of traffic means you have to make sure the results you're seeing are real. Once you've done that, though, you sure have a lot of information about latency and packet loss that isn't available without PingPlotter.

[Our knowledgebase](#) has a number of articles about certain conditions (ie: bad packet loss and latency at intermediate hops) and often stresses the importance of looking at the final destination first, then correlating that with the performance characteristics you're seeing in your network (with whatever other applications you use - web, VoIP, etc).

Network administrators rely on traceroute (the underlying technology used by PingPlotter) to troubleshoot networks. Because of this, most networks report predictably to the traffic that PingPlotter sends. Those network administrators who immediately dismiss the data are often those who are trying to avoid dealing with an issue. We've heard this objection since the first day PingPlotter was created - and in the meantime, our users (a huge number of them being network administrators) have been using it to identify and solve network problems.

PingPlotter is not a tool that gives you absolute answers in every situation. You do need to take your knowledge of the situation and apply your wisdom to what you see in the graphs. The alternatives to PingPlotter (and traceroute in general), for the layman, is...um...asking your network administrator (or your DSL/Cable provider, or your web hosting service) why the network is running slow. PingPlotter adds a level of information and precision that can be incredibly revealing.