Traceroute vs PingPlotter

Saved From: https://www.pingman.com/kb/article/traceroute-vs-pingplotter-160.html

Question

My service provider is refuting my PingPlotter data using Traceroute. What do I do?

Solution

Before we can understand why PingPlotter shows something different than Traceroute, it's helpful to understand how both technologies work.

Traceroute (or tracert) is a command line tool that is used to map out the network and get an idea of how long it takes to send and receive data to specific endpoints. Network technicians use this to determine what could be causing a network resource to be failing. It helps answer questions like: why can't I connect to the VPN? Why is Netflix buffering? Why is the Internet so slow?

When you start a traceroute, your machine will create a data packet and then send that packet out to the network. This data packet contains a field called Time to Live, or TTL. TTL specifies the number of devices, or hops, the packet can travel through before being tossed out. For example, if my data packet has a TTL of 30, it can only travel through 30 hops. If it gets to hop 31, then that device will send that packet to the data graveyard.

What does TTL have to do with traceroute?

Well, basically everything. Whenever a packet reaches a hop, the TTL value is reduced by 1 until it reaches the final destination, or TTL=0. Traceroute takes advantage of this in order to discover all of the devices between the computer sending the packets and the final destination.

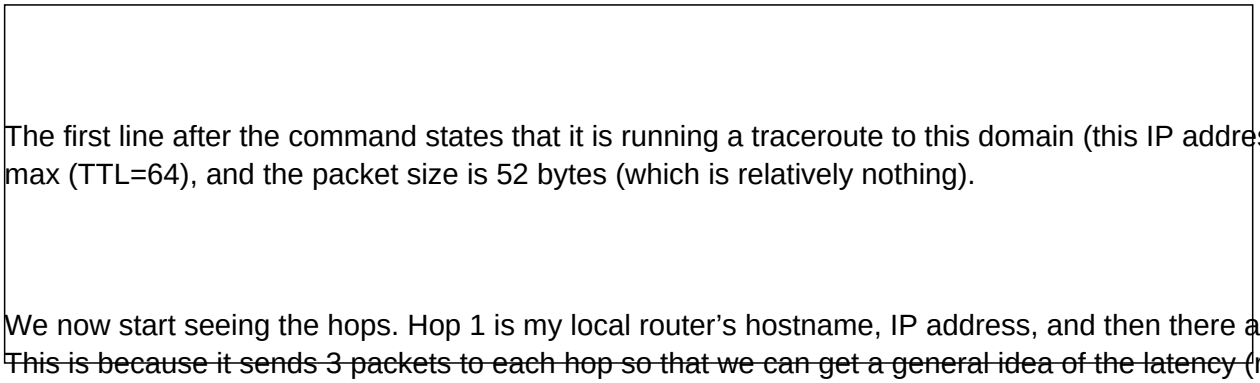
The first packet that traceroute sends out has a TTL of 1, which means once it reaches the first hop (typically your local router), TTL will be set to 0. The router will then send a message back to your computer to say, "Hey, you're out of TTL. You'll need a whole lot more TTL if you want to get past me, your router!" Your computer then notes the device it reached (your router) and the time it took to receive that response.

The next packet your computer sends out will have a TTL of 2, and then 3, and then 4, and then….well, you

This knowledgebase article and content copyright by Pingman Tools, LLC

get the idea. This goes on until your computer receives a response from the intended target, after which you should have a pretty good map of the network and an idea of how long each hop took to respond.

Let’s look at this in practice. Here is a traceroute from my Mac to pingplotter.com.



The first line after the command states that it is running a traceroute to this domain (this IP address), 64 hops max (TTL=64), and the packet size is 52 bytes (which is relatively nothing).

We now start seeing the hops. Hop 1 is my local router’s hostname, IP address, and then there are 3 times. This is because it sends 3 packets to each hop so that we can get a general idea of the latency (response time).

Hops 5, 6, 7, and 9 look a little different though. All we got back were three asterisks (***). What does that mean?

Anytime you see an asterisk, it means that no response was received. This could be for a couple of reasons, but the most common one is that the device at TTL=5 deprioritized the packet, or is configured to not respond at all to the particular packet type I sent.

Linux and Mac machines default to use UDP packets when running traceroute. Windows machines use ICMP packets for tracert. It might be that if I change my packet type from UDP to ICMP, I’d get a response….maybe. The important thing to know here is that the asterisks don’t necessarily mean packet loss, or that the device is offline.

And now you know more than you ever thought you would, or ever wanted to know, about traceroute. But this is important stuff!

Now that you understand how traceroute works, we can talk about how PingPlotter works. When you first start a trace in PingPlotter, it initiates a traceroute so that it can find all of the hops between the computer and the intended target. Once it finds all of those hops, PingPlotter starts pinging each of those hops every 2.5 seconds. Every so often it will run another traceroute to see if the route to the final destination has changed.

This knowledgebase article and content copyright by Pingman Tools, LLC

The reason PingPlotter pings each device instead of just running a traceroute each time is because the data in the packet is a little different than a traceroute. Ping's goal is simply to find out if something at the destination is alive, and approximately how long it took for the destination to respond. Ping sends two message types: type 8 (Echo Request) and type 0 (Echo Reply). When a ping is started, the sender creates an ICMP Echo Request packet to the destination. *If the destination is allowed to respond, it then replies with an ICMP Echo Reply. Please keep in mind that if no response is received, it does not necessarily mean the destination is offline or there is packet loss. ICMP traffic is very low on the internet totem pole and some devices simply ignore them. (Remember the asterisks from traceroute?)*

Let's look at a PingPlotter trace from my Mac to PingPlotter.com.

Notice how PingPlotter shows the same number of hops as on my traceroute, but I'm getting a lot more detail. This is because PingPlotter is constantly recalculating (and reporting/logging) the latency reported by the ping. This lets us track the latency over time. Keeping track of this latency over an extended period of time is so important because let's be honest, nothing is ever broken when the technician looks at it. Not to mention that network conditions are always changing.

And since PingPlotter is pinging each hop every 2.5 seconds, it has so much more data than traceroute could ever give you. I guess technically you could do it by hand but let's look at the numbers.

My trace to pingplotter.com has 10 hops. With a ping every 2.5 seconds to each one of those hops that puts us at 10 pings every 2.5 seconds which is 240 pings every minute or 14,400 pings every hour! That's a lot of data points! Good luck getting all of that by hand!

This knowledgebase article and content copyright by Pingman Tools, LLC