



Product Manual

Version 4.12.0

PingPlotter

(c) 1998, 2015 Pingman Tools LLC

All rights reserved. No parts of this work may be reproduced in any form or by any means - graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems - without the written permission of the publisher.

Products that are referred to in this document may be either trademarks and/or registered trademarks of the respective owners. The publisher and the author make no claim to these trademarks.

While every precaution has been taken in the preparation of this document, the publisher and the author assume no responsibility for errors or omissions, or for damages resulting from the use of information contained in this document or from the use of programs and source code that may accompany it. In no event shall the publisher and the author be liable for any loss of profit or any other commercial damage caused or alleged to have been caused directly or indirectly by this document.

Printed: March 2016

Publisher

Pingman Tools, LLC

Managing Editor

Pete Ness

Technical Editor

Gary Welch

Special thanks to:

All the people who contributed to this documentation and to the PingPlotter software.

Special recognition to:

Pierre Schwab

Tom Larrow

Jenna

Clara

Conrad

Andrea, D.D., Nina

Mele

Cam M.

Larry M.

Table of Contents

Part I Introduction	2
Part II 90 Second Overview	5
Part III How PingPlotter Works	12
Part IV Operation	16
1 Downloading and Installing	16
2 Basic Settings	16
3 The Interface - Graphs	17
4 Timeline Graphing	21
5 IPv6 and IPv4	24
6 Reporting	25
7 Route Changes	26
8 Uninstalling	28
Part V Common Tasks	33
1 Getting notified of a network problem (alerts)	33
Creating / Configuring Alerts	33
Troubleshooting Alerts	41
Email Setup (for alerts)	42
Alert Events	44
Event Notification	44
Event - Play a Sound	44
Event - Send Email	45
Event - Launch an executable	47
Event - Log to file	47
Event - Tray Icon Change	48
Event - Web REST Call	49
Event - Add/Remove from Summary	50
Event - Add Comment	50
Variable Substitution	51
2 Exporting data for further analysis	51
3 Long term monitoring and auto-saving data	53
Part VI Understanding Output	57
1 Discovering a route between you and a target	57
2 Interpreting Results - A Quick Example	58
3 Finding the source of the problem	60
4 Finding the source of the problem - part 2	63
5 Interpreting Results - Longterm Monitoring	65

6	Interpreting Results - Gamers	67
7	Interpreting Results - ISP Problems	70
8	Interpreting Results - Bad Hardware	75
9	Interpreting Results - Bandwidth Saturation	80
10	Voice over IP (VoIP) troubleshooting	83
11	Building a Compelling Case	95
Part VII Advanced Features		98
1	Raw TCP sockets and WinPcap	98
2	Command Line Arguments	99
3	Custom time graph intervals	100
4	Installer Options and MSI	101
5	Whois Setup	102
6	Running from a USB drive	103
Part VIII Reference		106
1	Options	106
	General Options	107
	Internet Options	110
	Email Settings	111
	Display Options	112
	Engine Options	115
	Auto-Save Options	119
	Route Change Options	121
2	Menus	122
	File Menu	122
	Edit Menu	123
	View Menu	124
Part IX More Information		127
1	Support	127
2	Purchasing	127
Index		129

Part



Introduction

1 Introduction

Your network connection probably isn't something you think about on a regular basis. When it isn't working right (or stops working all-together)... well, that's a completely different story. At that point, it's *all* you can think about. The only thing that's worse than your network connection not working? Trying to figure out why it's **not** working.

If you're reading this manual, odds are you've encountered (or are currently encountering) this exact situation. Or maybe you're just preparing yourself ahead of time (go you!). No matter your current network situation, PingPlotter can help you get to the bottom of these kinds of problems faster, so you can find a solution and get back to that blissful state of not having to think about your connection.

PingPlotter was originally created in 1998 to troubleshoot unacceptable lag in an online game (a problem which that particular ISP was claiming no responsibility for at the time). Over the years, the program has grown, and added a variety of features and capabilities. Today, it's a very powerful network monitoring, troubleshooting, and diagnostic tool, which is used by a variety of users - from the "weekend troubleshooter" to full-time network administrators.

PingPlotter can help with a variety of different network related woes - and can be a great help to you if:

- You rely on a network or internet service (like most of us do), which happens to be having problems - such as slow performance, random disconnects, or other similar issues.
- You're a systems administrator - and you need to know when connectivity to one of your servers go down (and want some evidence of where/when/why it went down).
- A provider is telling you that they can't see any problems (when you're clearly having issues) - and you need to show them where the problem really is.
- In general, if you're a user of something that relies heavily on a network or the internet, such as:
 - A web browser
 - VoIP services/video chat
 - Online gaming
 - Streaming audio/video
 - An ASP for your business (such as payroll, accounting, human resources, etc)
 - Home automation products


If you're here, one of the above bullets almost certainly applies to you (and none of them do... well, then

we're a bit confused as to how you got to this manual without a web browser). There are, however, a few situations where PingPlotter might not be the right tool for you:

- You have hundreds (or thousands) of network nodes with many services you need to monitor
- You need true SNMP capability (though, you can trigger SNMP traps via PingPlotter alerts)
- You need auto-discovery of network nodes
- You actually enjoy having network trouble on a regular basis, and would rather not know when/where problems arise

If you ever have any questions or comments concerning PingPlotter, this guide, or if you just want to email someone to say "hello" - please feel free to send us an email at info@pingplotter.com. We're always happy to answer any questions or provide any advice that we can!

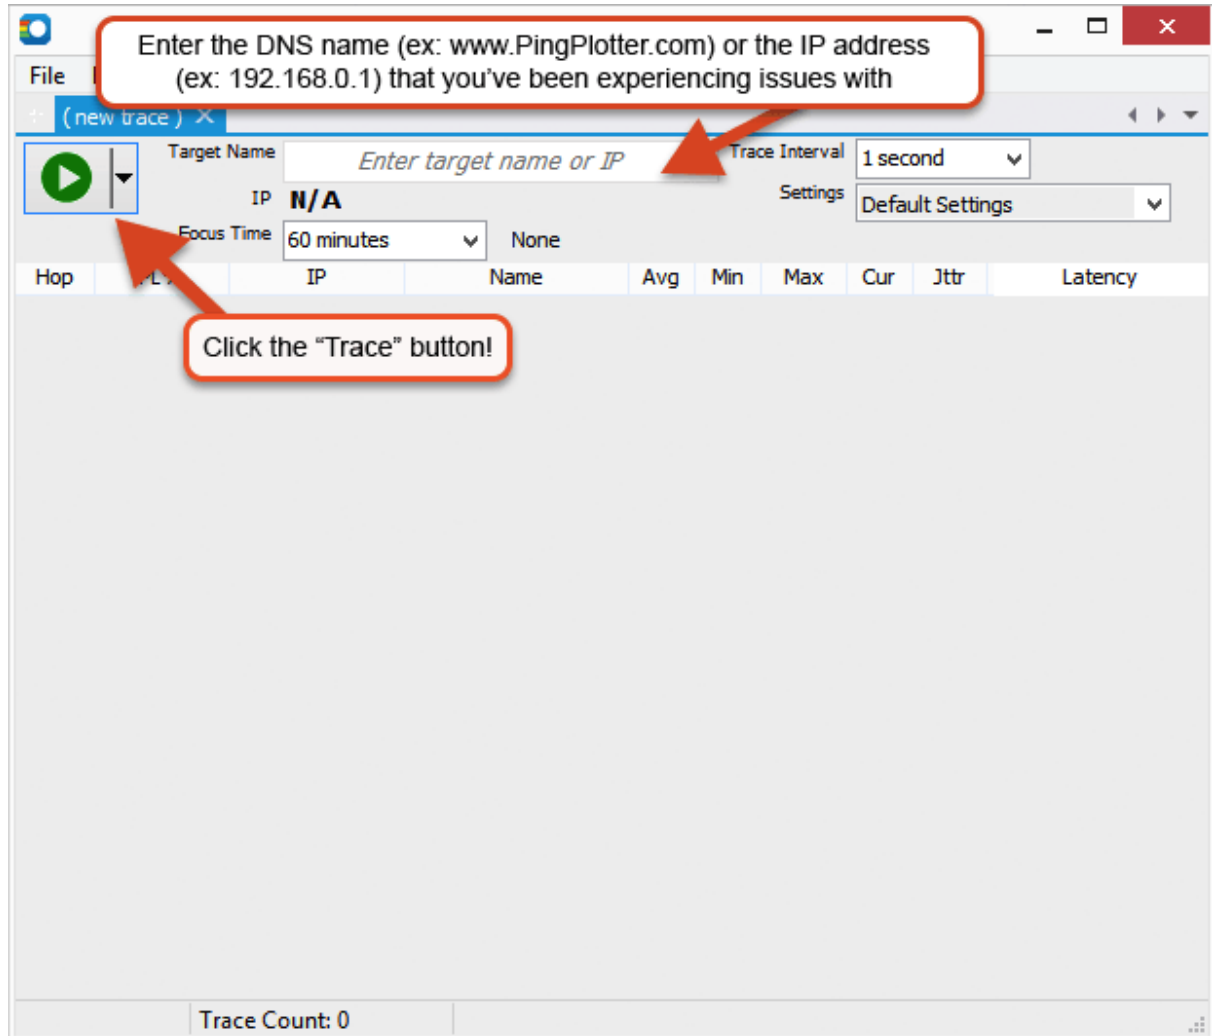
Part



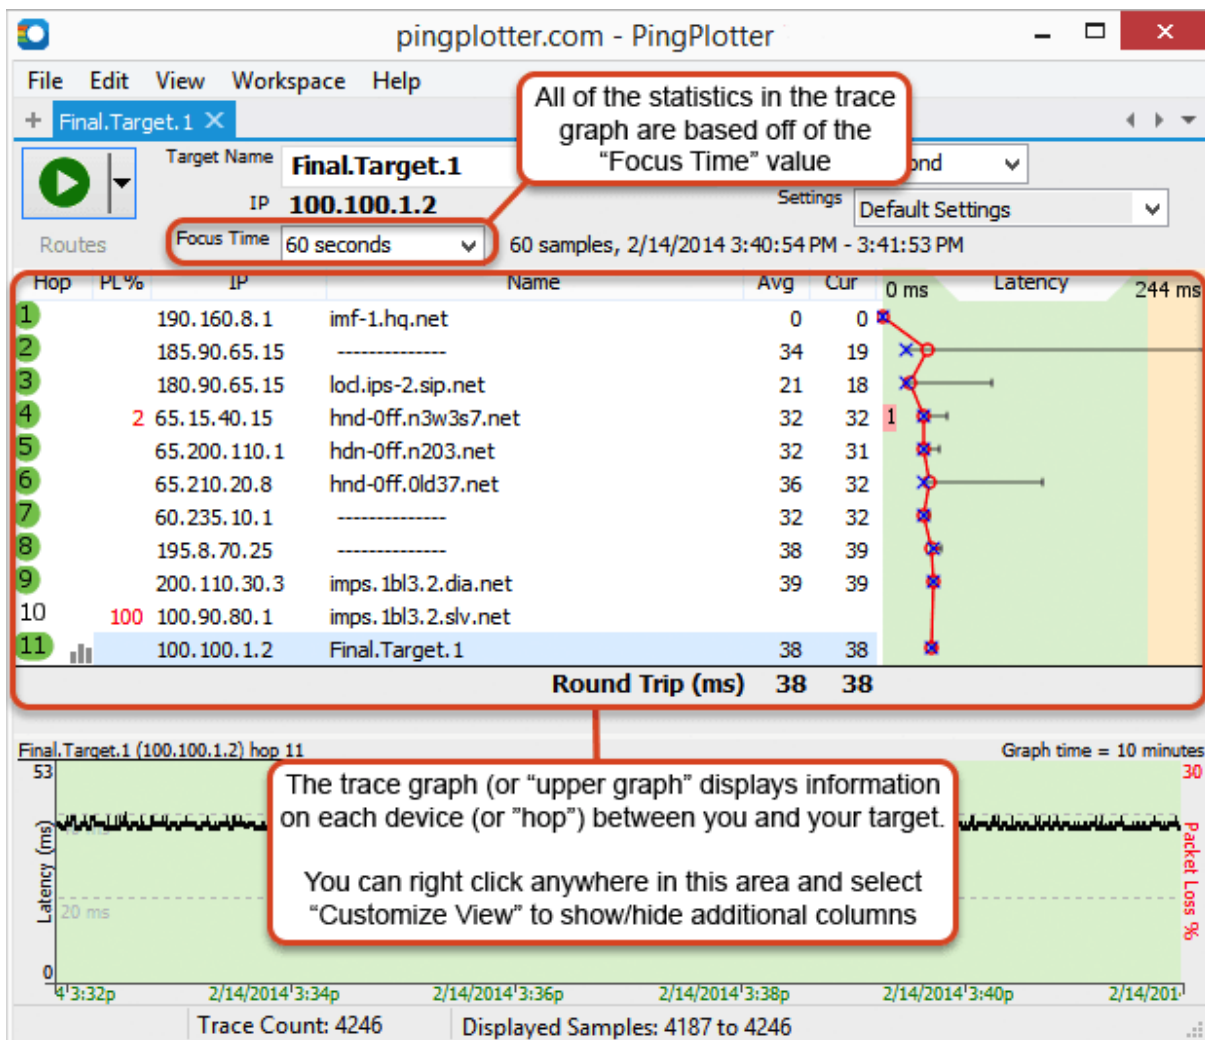
90 Second Overview

2 A 90 Second Overview of PingPlotter

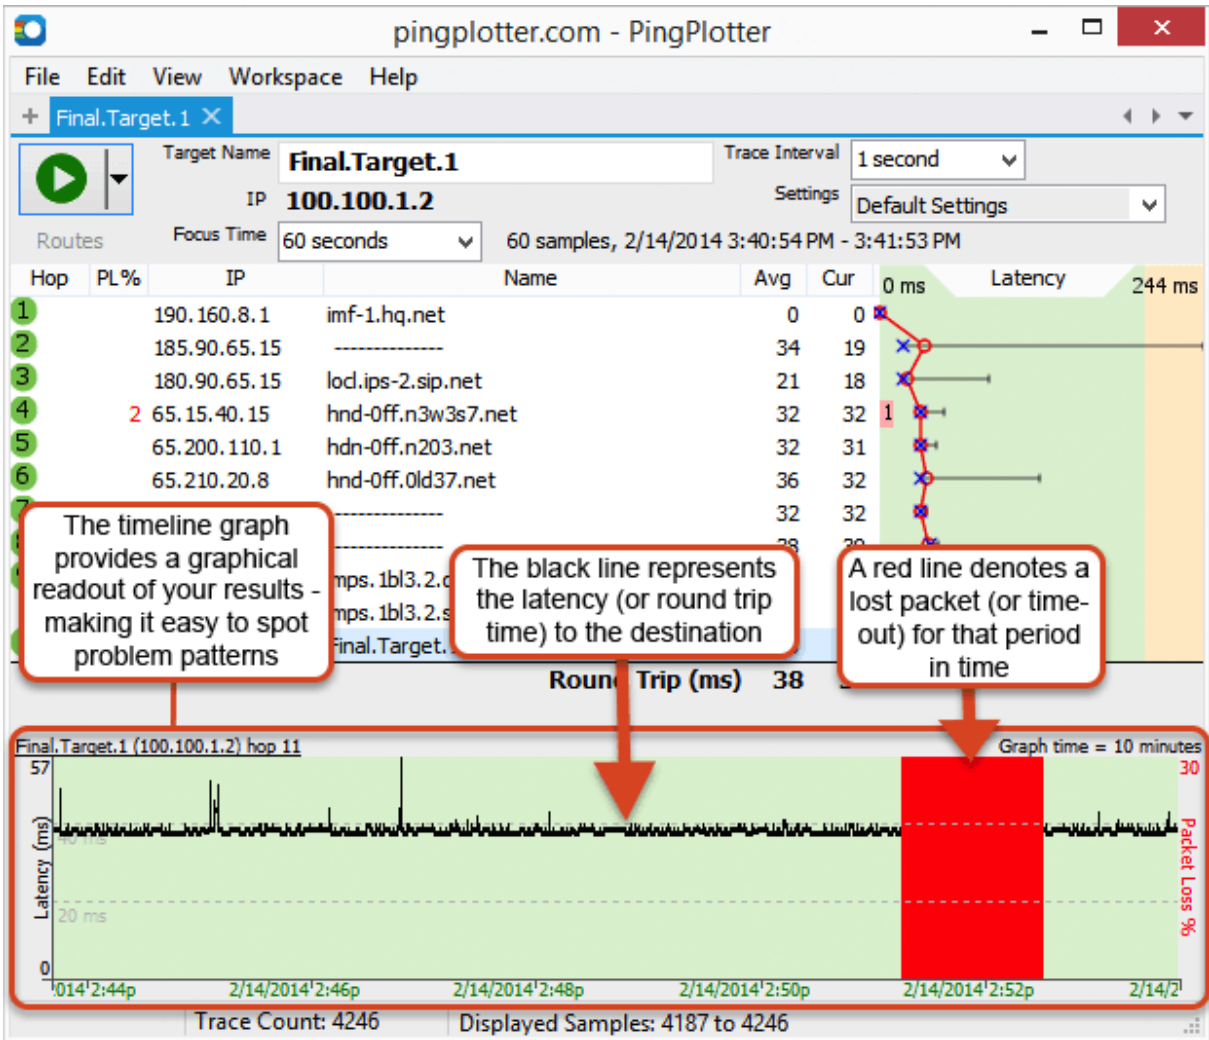
Once you've got PingPlotter [downloaded and installed](#)^[16], starting to collect data is a breeze.

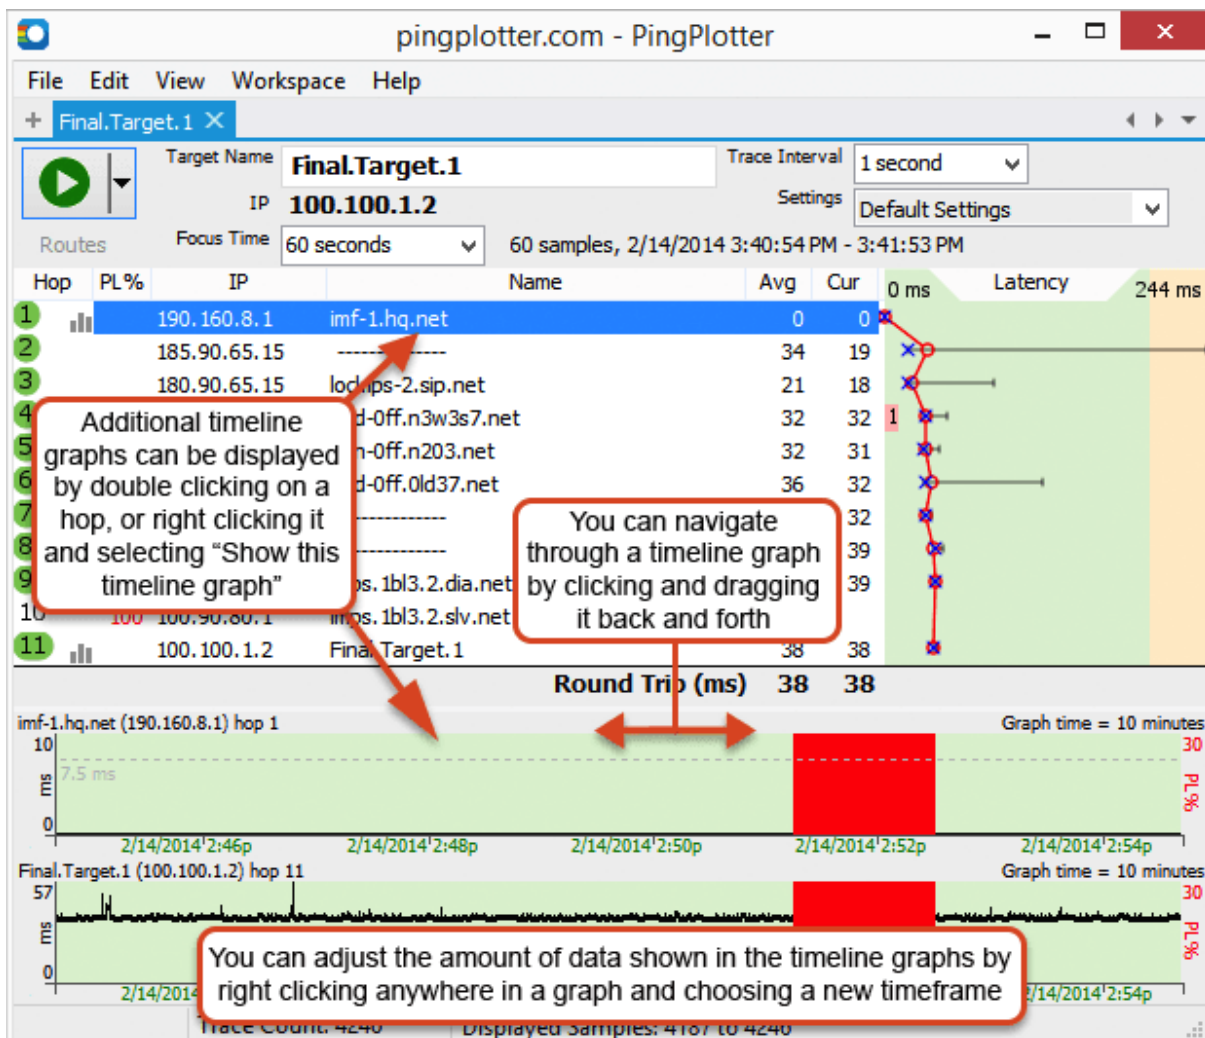


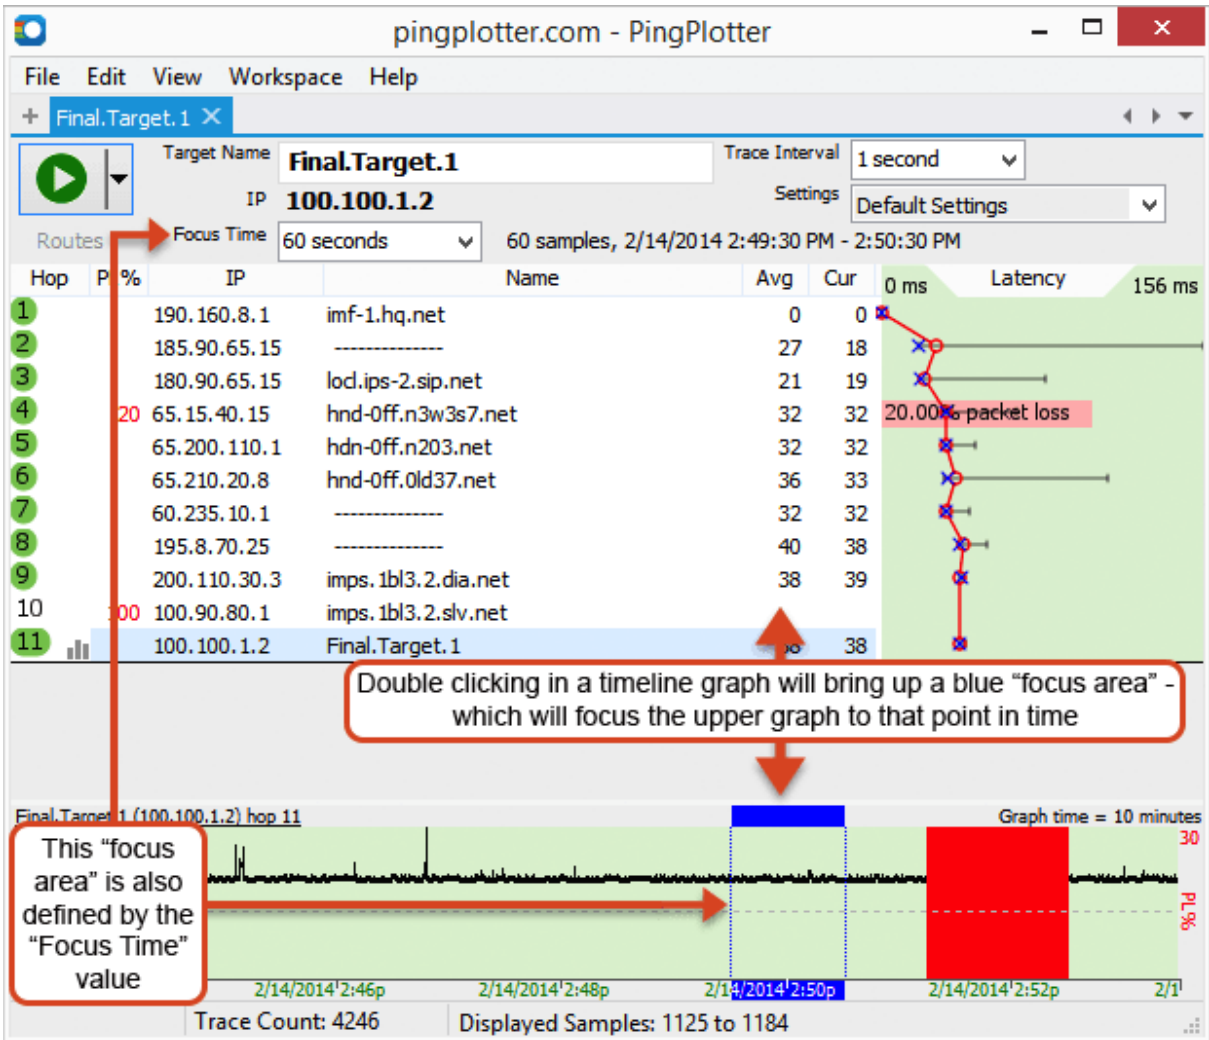
PingPlotter will then begin to generate route information. If you're getting a "Destination address unreachable" message - [have a look here](#) for some possible solutions.



Worried about the packet loss that you see at hop #4 and #10 in the above screen shot? As long as it doesn't seem to be affecting our final destination (which it isn't here) - then it's nothing to worry about. See this [Knowledge Base article](#) for more information.







The screenshot shows the PingPlotter Pro application window. The 'File' menu is open, displaying options such as 'New Target', 'Save Sample Set...', 'Save Image...', 'Export to Text File...', and 'Send Email'. Three callout boxes provide additional information:

- A callout box pointing to 'Save Sample Set...' states: "You can save your results as a ".pp2" file, and share them with others who may also be using PingPlotter".
- A callout box pointing to 'Save Image...' and 'Export to Text File...' states: "There are also options to save your results as an image or text file (if you're sharing with non-PingPlotter users)".
- A callout box pointing to the timeline graphs states: "***Tip*** If you're going to be sharing your results as an image - it's a good practice to have multiple timeline graphs open (as its helpful in identifying where any problems may be starting)".

The main interface displays a table of targets and their ping statistics, along with a graph showing ping times over time. The table data is as follows:

Target	IP	Host	RTT (ms)	Loss (%)
7	60.235.10.1		32	31
8	195.8.70.25		36	32
9	200.110.30.3	imps.1bl3.2.dia.net	38	39
10	100.90.80.1	imps.1bl3.2.slv.net	39	39
11	100.100.1.2	Final.Target.1	38	38

The graph below the table shows ping times for three hops: 'hop 1' (imf-1.hq.net), 'hop 5' (hdn-off.n203.net), and 'hop 11' (Final.Target.1). The x-axis represents time from 2:50p to 3:20p on 2/14/2014. The y-axis represents ping time in milliseconds. Red vertical bars indicate periods of high ping times or packet loss.

****Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details****

Part



How PingPlotter Works

3 How PingPlotter Works

If you're relatively new to the world of troubleshooting network problems, a lot of this subject matter can seem a bit daunting at first. Fret not, though - we've got an analogy that should help make things "click."

If you're not new to this space (and possibly just rolled your eyes at the idea of reading though an analogy), feel free to skip to the technical details toward the bottom of this section.

Network traffic is a bit like freeway traffic

A network can operate a bit like a freeway; things are great when everyone is going the speed limit and we only have 50% of the maximum traffic that's designed to go on that freeway. When we start to add more and more traffic, though, at some point the freeway won't have the capacity for any more.

Problems will start to arise as new drivers try to merge on the freeway. People that are already on the freeway will sometimes slow down and cause traffic jams. If it gets too bad, some people may give up all together - and decide not to continue their journey.

On a freeway, this might be referred to as "congestion" - and this "congestion" happens on networks as well (in pretty much the same way). Packet loss and latency (two terms you'll be getting very familiar with) are both symptoms of congestion - when there's too much traffic for the network to handle.

Lets say this proverbial freeway is the one we take to get home from work everyday. Let's also say (for this example) that between work and home there are 15 off-ramps with turnaround points off of the freeway. Lets also say that we've got a team of 15 people, each with their own car, ready to do whatever we ask of them (again, for the sake of this example. Stick with us here).

If we want to find out the conditions on the freeway for our drive home from work, we could send out 15 cars, and assign each driver to one of those 15 off-ramps. The instructions for each driver are the same: get you your assigned off-ramp, turn around, and come back. Then we'll measure the time it takes each car to get from us, to their off-ramp, and then back to us.

The most important car is the one that goes all the way to your target (or home) - that 15th car. If it makes it there and back again in the expected time, then we know that the traffic on the freeway is running pretty well.

If that 15th car takes longer than expected (or if it never returns), then we can look at the results for the other off-ramps to find a likely place where problems could be occurring. Maybe all the cars through off-ramp #9 had no problems (and returned quickly) - but the cars that went through off-ramp #10 (and beyond) started getting delays (or didn't return). From this, we can see that there is some kind of

problem happening past off-ramp #9.

PingPlotter operates *very* similar to this; it sends out data packets that go all the way to a target destination (as well as each stop in between), and measures the amount of time it takes to get there. It also measures how often a packet (or a router) gives up. This information can be used to figure out where (and when) problems are occurring.

For the sake of taking this analogy too far, lets look at one other scenario. Let's say that off-ramp #5 is in a small town where the police are are of a disposition to pull people over for no reason (this is all theoretical, of course - we would never imply that any police actually do this). Each car that we send to off-ramp #5 have to pass buy these police officers, and 20% of the time, they get pulled over. Another 15% of the time, there's someone else pulled over there, and our car has to wait while that car moves off the road. Meanwhile, traffic is whizzing by on the freeway, unrestricted.

This situation can happen on a network with PingPlotter as well - where the packets going to hop 5 might get waylaid by some local rules and show packet loss, latency, and jitter that are not being experienced by packets destined for other places.

Technically speaking...

At its heart, PingPlotter is a trace route utility. It's souped up and on steroids, but the basic data it collects is based on the theory of trace route.

A ping packet is an IP packet requesting that a copy of its contents be echoed back to the sender. When you "ping" a site, you send over an echo request and that site responds back that it received it.

One of the parameters on a ping packet is something called "Time to live" (TTL) - which is an IP header field designed to keep packets from running in loops (essentially forever) throughout a network (this can happen when there is a route change, and the routers involved don't all know the same information as new information is being replicated out). Initially it's usually set to somewhere between 64 and 255, and is reduced by 1 every time it passes through a server. If the TTL should ever reach zero, the packet has expired, and the router that it's passing through will send it back to the source.

Trace route plays with this TTL number on outgoing packets. It first sends out a packet with a TTL of 1. The first router that sees this and decrements it to 0, and then sends it back. It also sends back its own IP address with the packet, and DNS is used to do a lookup for an actual domain name.

Next, traceroute sends out a packet with a TTL of 2 so it can find out what the next computer in the route is. Then it sends out a packet with a TTL of 3. This process is repeated until the final destination is reached. At that point, you know the entire path the packet has traversed to reach the destination

computer/router. Each server/router in this chain is called a hop.

This method can help us determine the route a packet takes, but if we time each of these packets, we also know how long it takes for a packet to make it from our source PC, to that router, and then back again. This is called latency.

The last hop in a (successful) trace route is actually the round-trip time to the destination server. This is an important concept to understand. You don't add up all the times between you and the destination host - as that time has already been added. The time to the last hop in the chain is exactly the same as is if you'd used a ping utility to that host. So a trace route utility is actually two utilities - ping AND trace route.

PingPlotter speeds up this process by sending out packets to the first 35 servers in the route all at the same time. This makes a HUGE difference in overall speed. It also means that the network conditions for each hop are very similar - so the numbers are better compared.

Part



Operation

4 Operation

4.1 Downloading and Installing

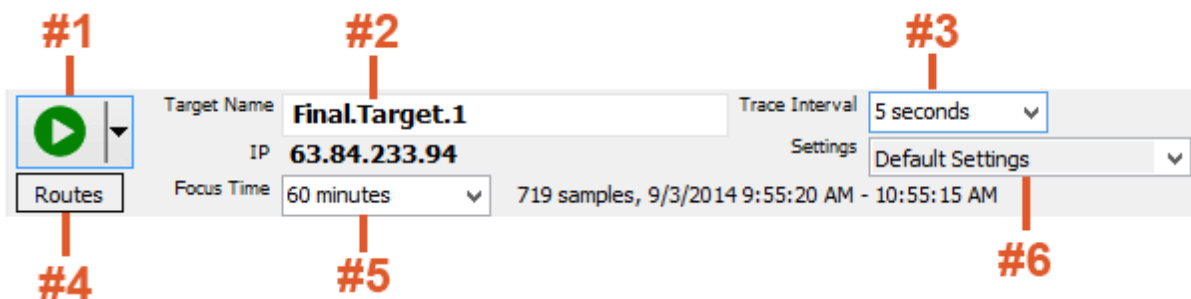
- Open <http://www.pingplotter.com/download.html> in your browser.
- Click on "Download" button. If you have problems downloading via HTTP, you can alternately click on "Download via FTP" option
- When asked by your browser what you'd like to do, choose "Save"
- Navigate to where you saved the PingPlotter installer, and double click on it

The PingPlotter installation program then starts the wizard driven install that steps through the installation. We suggest leaving all the settings at their defaults.

At the end of the install, you'll be asked if you want to "Run PingPlotter now". If for some reason you don't want to at that time, uncheck that checkbox. Otherwise congratulations, you now should have a PingPlotter group in with the rest of your installed applications (i.e.: Start menu, Programs group, PingPlotter folder) and are ready to go!

If this procedure doesn't work for you for some reason (if you get a weird error, etc.), please see our knowledge base article [here](#). If you still encounter problems with the download, please send an email to support@pingplotter.com describing the specific error you're getting, and/or any symptoms that you're seeing.

4.2 Basic Settings



1 - **The Start/Pause button** - clicking the drop-down menu on this button will provide options to reset and restart your trace, limit your trace count, or trace continuously. If you're using PingPlotter Pro, you can also start tracing to a new target, or create a new summary screen from here

2 - **The Target Name** field is where you enter the IP Address or DNS name of the destination you want to trace. If you haven't entered anything in this field, clicking on it will bring up a "history pane" that

shows targets you've recently traced to. If you enter an IP Address here (and press the start" button, or the enter key on your keyboard) PingPlotter will start tracing immediately before the IP is resolved to a name. The name will show as "resolving" until the request is complete.

3 - **The Trace Interval** is the amount of time PingPlotter will wait between each sample set. If you're doing a long term monitoring project, you may want to set it to be longer (15 - 60 seconds). If you're doing a quick test, you might want to set this to something lower (1 - 5 seconds). If the up/down arrow doesn't have the amount of time you want, just type the time interval you want (e.g. 3.5 seconds).

4 - **The Route Change** button is used to show the history of route changes. Anytime any hop in the route changes, PingPlotter stores the old and new route data and adds the time of the change to this list box. Double-clicking on any time will show the route as of that time. This is the starting time for the change. If the route change button is grayed out, then there is no new route change information. If it has a black border (as seen in the about screenshot), then new route change information exists.

Double-clicking on the time-line graph (covered in the section: "[The Interface - Graphs](#)"¹⁷) will refocus the upper graph to the period you double-clicked on the lower time-line graph. The route window will also follow this - to show the route that was current at the time you selected.

5 - **The Focus Time** dictates how much recent data PingPlotter will use to calculate its statistics (All numbers in the trace (upper) graph are affected by this). This field can be set to look at a time (5 minutes, 1 hour, etc), or a certain amount of samples (10 samples, 1000 samples, etc). When running a trace, PingPlotter can look at just the most current samples. This is great to watch "trending" (where the response changes over time). If you include ALL samples (type 0 in this field for ALL), then after a lot of samples, new samples don't affect the graph very much. Setting this to something like "10 samples" allows you to see how the response times are right now.

6 - **The Settings** option (PingPlotter Pro only) allows you to toggle through any named configurations that you may have set up (so you're able to utilize different settings on different targets)

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

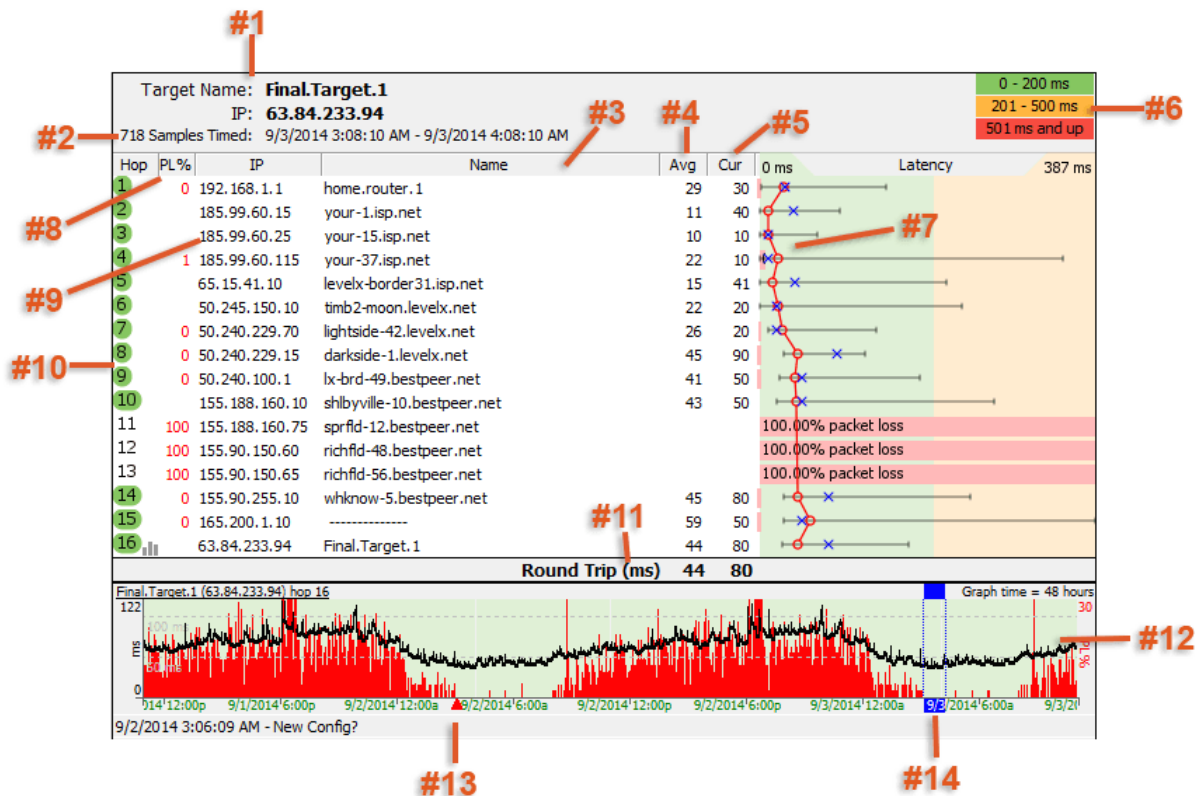
4.3 The Interface - Graphs

The graphs are where PingPlotter really shines. At a glance, you're able to visually see where a problem lies. There are actually two graphs available, the **Trace Data Graph**, and the **Timeline Graph**. We'll explore both in this section, as well as some other items related to them. Please refer to the below image (annotated with numbers) that we've saved from PingPlotter ("File" -> "Save Image"), and the

explanations (referencing the numbers) below the graphs.

The upper graph is called the **Trace Graph**. All columns on the trace graph are re-sizable. The lower horizontal graph is called **Timeline Graph**.

All numbers on the trace graph use the "[Focus Time](#)"¹⁶⁷ setting to control how many samples are used in the calculations. By default, you're looking at the most recently collected samples - but you can also focus on samples that are not the most recent samples by double-clicking on the time graph. This will focus the upper graph on that period you double-clicked (we cover this topic in a bit more detail below).



- 1 - The **DNS name and IP Address** for the host you're doing the trace to - in nice big letters so you quickly know which trace you're looking at.
- 2 - **Beginning and ending times/dates** for the trace. Very useful if you're saving off graphs. It's nice to know the time window the trace was done in. On "Copy as image" or "Save as image"-created images, you'll see the "XXXX samples timed:" value used to create this graph (where XXXX = "Focus Time"). In the live application, this line isn't shown because it is available elsewhere.
- 3 - This column shows the **DNS name** of the device for that hop. A "-----" in this column indicates that PingPlotter was unable to resolve a name for that device's IP address. This is not a flaw in PingPlotter, it just means that your DNS server doesn't have a name for that IP address (or that

address just doesn't have a DNS name, period).

- 4 - The **Avg** column shows the average response time of the last "X" samples (where "X" is the "Focus Time"). Any timeouts/lost packets are not included in this value.
- 5 - The **Cur** column shows the individual sample time of the most recent sample included in the set. If a number is displayed as **ERR**, that means the packet was lost (i.e.: a packet was sent out, but never made it back).
- 6 - Shows you the thresholds you've set for the colors on the trace graph background. In this case we've set 200ms as the warning color, and 500ms as what we thought was a critical speed (PingPlotter defaults to 200ms and 500ms for these values, however, you can change them [in the Display section of the Options dialog available under the Edit menu](#)^[112]).
- 7 - The red line on the graph represents the average response time for each host for the currently selected samples.
 - The **blue X** represents the response time for the current packet (this can be turned off if you're sending a copy of your graph to somebody - some folks find it confusing when they're not actually there watching the live trace).
 - The **black horizontal lines** represent the minimum and maximum response times.
 - The **red horizontal bar** shows the packet loss for that hop (same as the PL% column, but there for readability). PingPlotter uses a dynamic scale for its graph. The bottom number is usually 0, and the top number represents the maximum response time in milliseconds. If you wish, you can change this to a fixed scale [in the Display section of the Options dialog available under the Edit menu](#)^[112].
- 8 - The **PL%** indicates the number of packet(s) that have been lost in the current sample set. If you're only including the last 10 samples, then only the number of lost packets in the last 10 samples are shown here. If you want to find out how many time-outs have happened over the entire session, change the "Focus Time" to "0" or "All."
- 9 - The **IP address** of the router that reported back for that hop.
- 10 - **The number of hops** that device in the route is from your computer. If a hop has a bell icon next to it, this means that hop is being monitored for an alert ([alerts are covered in the advanced settings section of this tutorial](#)^[33]). Multiple alerts can be configured for the same IP, and alerts don't work unless some IP in your current route is being monitored. If a hop has a graph icon next to it (like hop #16), that hop is being traced on a time-line graph.
- 11 - The **Round Trip** line is basically there for ease of reading. It's the same value as the last server in

the route. This is the time it takes for a ping to get from your computer to the target device and back.

12 - The **Timeline Graph (TG)** is one of the most powerful features in PingPlotter, and great for long-term monitoring projects you may be doing.

- A red line on this graph denotes a time-out for that period. **Double clicking anywhere on the graph shows you the trace data detail for that period in the upper graph** (denoted by a blue "focus area" like you see in the example above - the "Focus Time" setting controls the number of samples encompassed by the "focus area"). This allows you to see what was happening along the route when that particular time-out occurred.
- Right clicking on the TG allows you to pick the time period you want to view - from 60 seconds up to 48 hours. This value affects all timeline graphs. If you'd like to add additional time periods you can do so by modifying your PingPlotter.exe.ini file. This is covered in the [Advanced Topics and Tips - Timeline Intervals](#)^[100] section of this tutorial.
- You can slide the graph (by doing a "Click-Drag"), allowing you to look at the past data for this trace. You can right-click and select "Reset focus to current" to move all graphs back to the current time. If you've been running the trace for a long time, it's helpful to adjust the time period you have set before you start moving back in the history, i.e. if you're going back a couple of days you might want to set the time period to 24 hours first.
- You can see a TG for any hop by right clicking on that hop in the trace graph and selecting "Show this timeline graph". You can also get rid of that graph by right clicking and deselecting it. A graph icon appears next the hop # to show that a graph is being shown for that hop. A shortcut for this is to double-click on a hop # to show / hide the time graph.
- We discuss time graphs a bit more in the "[Timeline Graphing](#)"^[21] section.

13 - Comments are denoted by a red up arrow, and are a very handy feature enabling you to add a comment (that is saved when you save the sample set) for things like planned outages, configuration changes or whatever else you want to make a note of.

14 - Double clicking anywhere on the timeline graph will bring up a Focus Area, which will focus the trace data graph (top graph) to that particular point in time. This is particularly useful for investigating spikes or time-outs


Tools and other options available for the Trace Graph

- You can display the Minimum and Maximum columns by right clicking on the upper trace graph and selecting them (via the "Customize View..." option).
- You can copy the IP address or DNS name for a hop to the clipboard by right-clicking on that hop, selecting the Clipboard option and then clicking on what you want to save.
- You can do a Whois on a particular hop by right clicking on it and selecting Whois Information. Note that by default this queries whois.crsnic.net. You can add additional Whois servers by editing your PingPlotter.exe.ini file. Instructions for doing this are in the [Advanced Topics and Tips - Whois](#) section.
- You can lookup who owns the particular IP range a hop is in by right clicking and then selecting "IP Block Lookup (ARIN)".

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

4.4 Timeline Graphing

Network problems can often happen when you're not watching for them. The timeline graph feature in PingPlotter gives us a quick way to look over a visual representation of our trace data. This makes spotting problems (or potential problems) much, much easier.



By default, PingPlotter will automatically trace the last hop (the host you're tracing to) on a timeline graph.

In PingPlotter Standard and Pro you can also display a timeline graph for any of the other hops in a route by either double clicking on that hop, or right clicking and selecting "Show this timeline graph." You can also turn off any graph by these same mechanisms.

The amount of data displayed on the graph can be changed, too - just right click anywhere on the graph

and select the amount of time you'd like to display (this will affect all timeline graphs and is saved when you shut down PingPlotter).

Navigation

If there is more data collected than we can show on a timeline graph (for example, if you've got 48 hours worth of data, but have your graph scale set to "10 minutes"), you can click (and hold down) your mouse button on the graph, and drag it back and forth. This allows you to move back in history and examine the samples during those times.

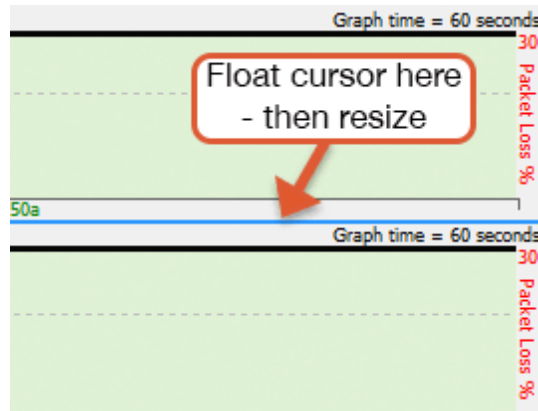
The scroll wheel on your mouse also works some magic with the timeline graphs. If you click on a graph and scroll down, you'll move back on the time graph, scrolling up will move the graph forward. If you hold down the "CTRL" key on your keyboard, scrolling up or down will toggle through the time period options on the graph (so you can basically zoom in, or zoom out using this method).

If you're thinking to yourself "it'd really be great if I could navigate the timeline graphs without having to reach over and touch my mouse..." - well, then, you're in luck! There are some keyboard shortcuts available that can help you move around:

ALT-HOME	scroll to the beginning of the collected data
ALT-END	scroll to the end of the collected data
ALT-LEFT	scroll back in time (about 5% of the graph width)
ALT-RIGHT	scroll forward in time
ALT-PGUP	scroll back in time (about 1/2 of the graph width)
ALT-PGDN	scroll forward in time

Resizing the time graphs


To re-size the time line graphs, scroll to the last viewable graph (the scroll bar should be at the bottom of its range). Now, float your mouse cursor between the lowest graph and the one above it, and you should get a resize pointer and you can then size the graphs at your pleasure. All time graphs will re-size to match.



Comments


If you're using PingPlotter for long-term monitoring, you may sometimes run across a situation where your network was effected by something you know about (power outage, big file download, that time you threw your router out the window, etc). You may also run into situations that you don't control, but know the cause of (or can speculate on). Being able to take notes about these situations and tie it to your data can prove to be very, very helpful.

To create a note, all you need to do is right click on a timeline graph (at a point you want to create a note), and select the "Create Comment" option. A prompt will appear asking you to create your comment, and then PingPlotter will draw a red triangle on the lower edge of the timeline graph. If you float your mouse over the triangle you can see the note:




Right clicking on the comment triangle will also allow you to edit or delete it.

When you save data as an image from PingPlotter, your comments will get attached to the image (along with the times they happened). If you're sending a image to a network provider, this can be especially valuable - as it helps explain the events on the data that you're sending to them:



The Focus Area

Any time you double click on a timeline graph, a blue “Focus Area” will appear, which focuses the upper graph to that point in time. This focus area is based off of the “Focus Time” value (and this won’t work if you have this value set to “0” or “ALL”).



If you’ve got your graphs set to show 48 hours worth of data, and if you find a time period that looks like it might be interesting, you can double click on the timeline graph at that point and the trace graph will move to that time period as well. You can then change your timeline graph scale, and the lower graphs will stay focused on the period you selected. This makes it easy to spot, and zoom in on problems.

When you’re finished going through your graph history, you can reset everything to display your current results by right clicking on a timeline graph and selecting “Reset focus to current.” This returns both of the graphs to, you guessed it, the current time.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

4.5 IPv6 and IPv4

Starting with v4, PingPlotter supports IPv6 capability for ICMP DLL packets.

PingPlotter will automatically pick the right address type when you enter a name. It uses the Windows-

standard name lookup utility, so it will pick an IPv6 address over IPv4 based on how Windows works. Some targets (a growing number of them, "www.google.com" is one example) will let you use their services either way.

If you know specifically which protocol you want to use, you can prefix the name with the IP type, like this:

```
ipv4:www.google.com
```

or

```
ipv6:www.google.com
```

If your computer supports that IP type, then PingPlotter will use it. If not, you'll get an error.

4.6 Reporting

PingPlotter doesn't have any *printing* options (who needs paper anyways?), but it does have a variety of options available to help output data (which will let you manipulate it in your favorite software package from there).

There are several methods here:

Raw data (in TraceRT format)

In PingPlotter, you can choose the "Edit" -> "Copy as Text" option to copy the raw data to your clipboard. This is copied in a format that's similar to *most* text-based trace route programs.

It's important to note that the "Focus Time" option on the main screen is used to decide how many samples to include in this. If you have the "Focus time" set to 10 samples or higher, the menu option will show as "Copy as Summary Text" instead of "Copy as Text." Holding down the shift key when you click on the "Edit" menu option will switch between modes.

The graph

There are a few ways get a picture of your PingPlotter graph(s). The quick way is to select the "Edit" -> "Copy as image option" (which will copy the graph you're viewing to your clipboard). The column and graph sizing will match up exactly to what you see on your screen - so make sure everything you want to show is explained.

Another method is the "File" -> "Save as Image" option. You can also have PingPlotter automatically create these images for you as well by scheduling an auto-save option.

Comma delimited text file

This option is built to import into a program like Microsoft Excel - so you can manipulate the numbers and create output in different formats. To export a comma delimited text file, use the “File” -> “Export to Text File” option. From there, you’ll be able to specify the file you want to export (and a couple of other options). You can either export all samples in memory, or the range as specified on the main screen.


The “Include sample times in export file” option will specify whether or not to include the time each sample was taken at. If you don’t have this turned on, all of the samples will be output, but you won’t get corresponding times. Enable this option to include the times.

We discuss more reporting options [in our knowledge base](#).

4.7 Route Changes

A route change happens when packets change the route that they take between you and your selected target (and can occur at any time). Route changes are a pretty normal fact of life on the internet - and can happen for a variety of reasons (load balancing, to route your data around a problem area, someone may have fixed a problem, etc)

PingPlotter keeps trace of all route changes that it discovers when tracing to a target. When PingPlotter sees that a route has changed, the “Routes” button (found on your trace graphs) will light up:



Clicking the “Routes” button will pull open a side panel that shows a list of each route change (with the most current being the one on top). To see what route the data was taking at any time, just click on that time in the route change window. The trace graph on the right will display the route being taken at that time.

If you’ve been surfing through your list of route changes for a while, and you notice that you’re no longer looking at the most recent route or time period, just right-click on a timeline graph and select the “Reset focus to current” option.

It can sometimes be hard to see what is different between two routes - but lucky for you, we’ve made comparing routes pretty painless. If you pick a route, and drag your mouse cursor to cover other routes that you’re interested in comparing (or use the old tried and true “CTRL” and “SHIFT” keys to multi-select

routes). The trace graph on the right will highlight the hops that changes. Once you see the hop that's different, you can toggle between routes by clicking on them.

If a route is changing a lot while you're doing this work, pausing PingPlotter can sometimes make things a bit easier to work with (just don't forget to un-pause when you're finished!).

Route Change Issues

Normally, PingPlotter is able to keep track of route changes pretty well, and pretty much without notice by anyone (unless you're specifically looking for it). The time when it can start to cause problems in the data that PingPlotter displays is when the length of the route changes (when your destination shows up at different hops - depending on the route being used) - as the changing routes cause problems in the final hop.

If you want to suppress recording information about route changes, you can do this in PingPlotter unless the length of the route is changing, in which case recording these changes can't be suppressed.

Now, there are a few different variations of things that could cause problems, and causing your packet loss. A big variable in this is whether or not your route length is changing.

One thing to know about PingPlotter (and all ping tools) compared to HTTP web access is that HTTP uses error correction in its communication. If you get a lost packet when transferring HTTP, you often don't notice this because the protocol corrects for errors. The ICMP protocol (which is used by PingPlotter and other ping tools) is lossy - so if something along the way drops data, it's never corrected, just reported by whatever tool sent it out. This is one possible reason why you're not seeing lost data when browsing the web or downloading something, but are seeing it with PingPlotter.

The symptoms seen when data is lost in an error correcting protocol is that performance suffers. When data is lost, the protocol negotiates for it to be resent and this takes time. If you're seeing slow performance when downloading files, or browsing the web, it's possible that the drop in performance is being caused by packet loss.

Creating a Route Change Mask

If you're getting a lot of route changes that are happening normally (like if there's a load balancing router in your route) you may want to "mask" these route changes.

If you right click on a route that has an oscillating router, you *should* get an option to "Add route change mask" to hide these route changes in the future.

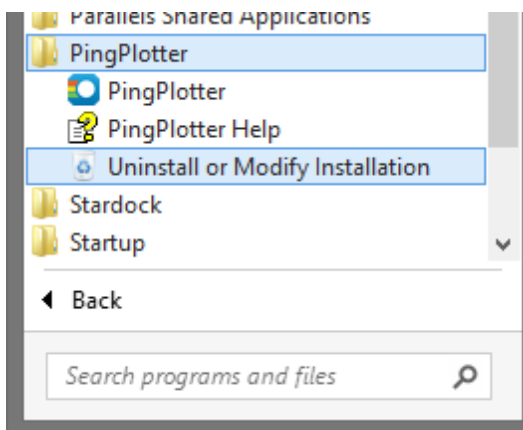
If that option *doesn't* appear, it may be because the IP addresses are too different, and PingPlotter

can't see that it's an oscillating router. If these ends up being the case, try multi-selecting the routes (in the route-change panel) where this oscillation is happening and try to add the mask again.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

4.8 Uninstalling

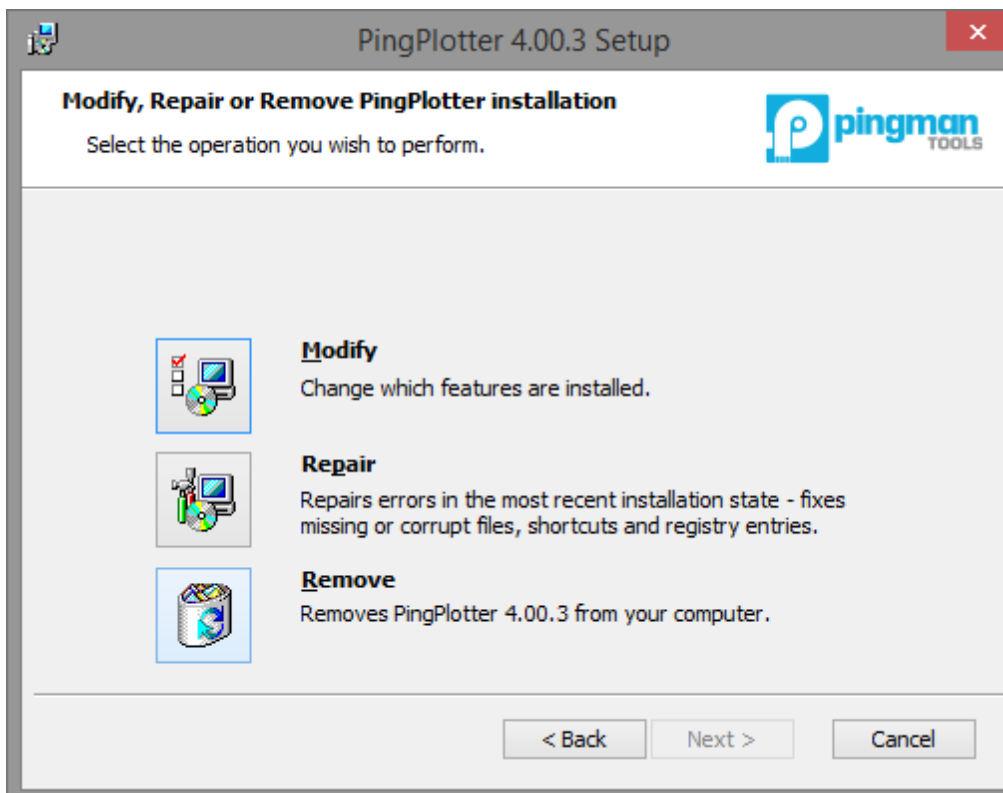
To uninstall PingPlotter - just locate your PingPlotter directory in the start menu and run "Uninstall or Modify Installation":



Click "Next":



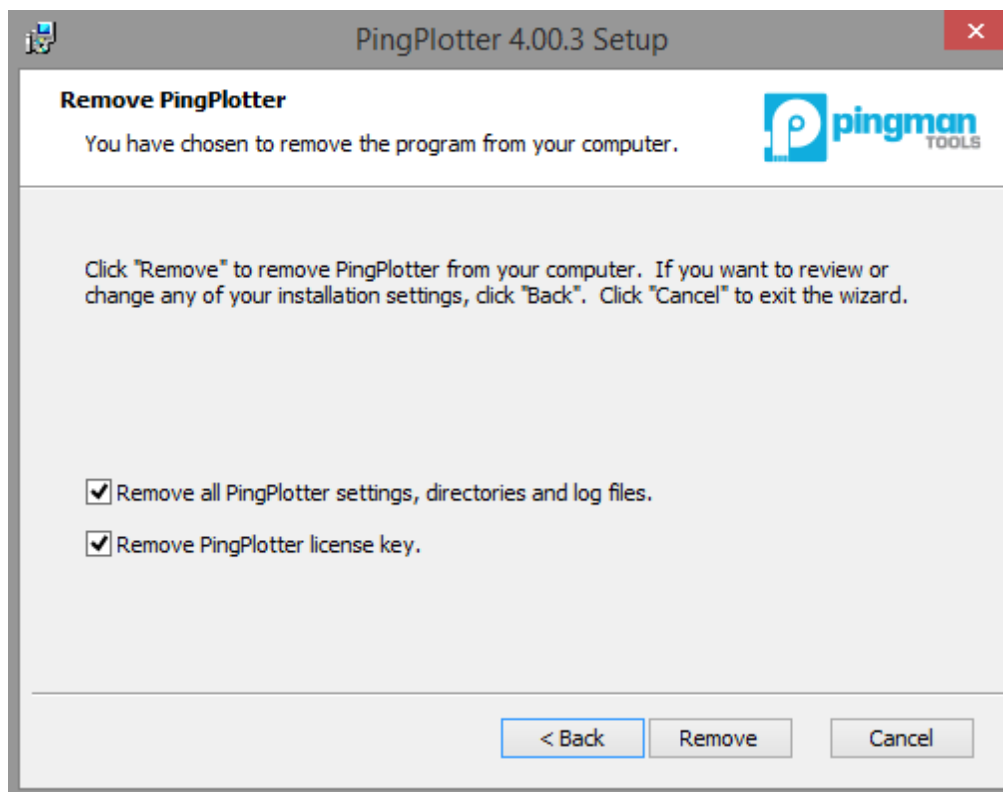
Click the "Remove" option:



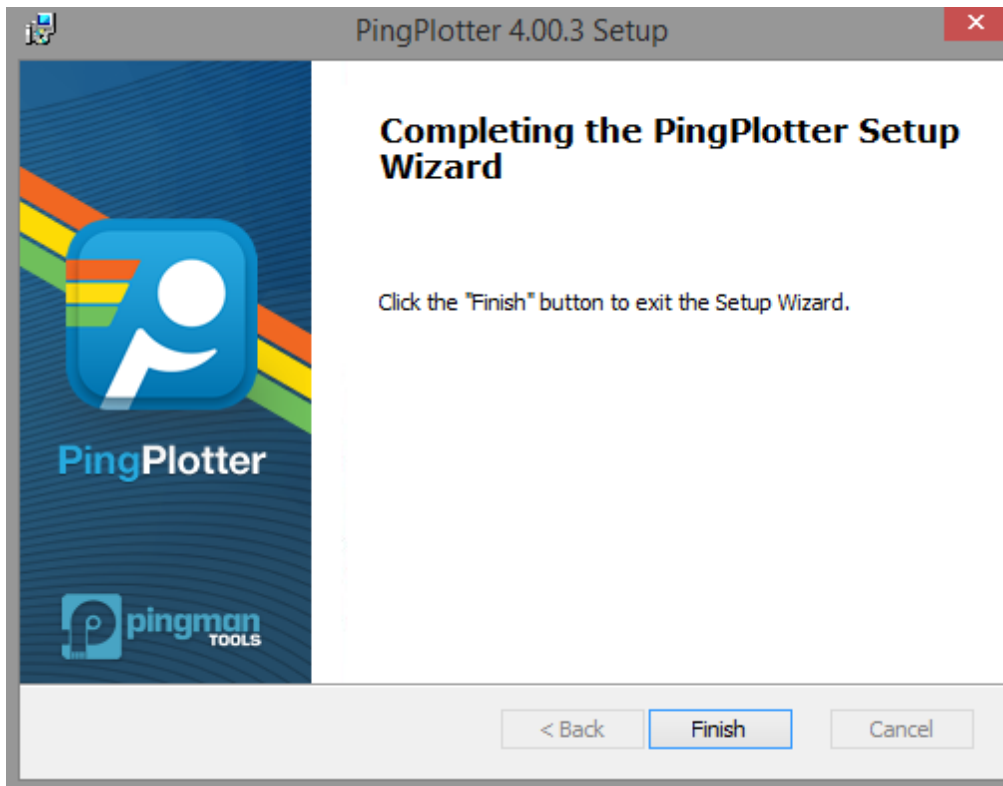
If you want to keep your license key, setting, directories, or log files, deselect the appropriate boxes.

Note: if you are uninstalling through Add/Remove programs these options are only available though the "Change" option. The "Uninstall" option does not provide these options.

After your preferred boxes are checked or unchecked, click "Remove":



Click finish to complete the uninstall:




You've successfully uninstalled PingPlotter!

Troubleshooting

If you can't find the Modify or Remove option in the PingPlotter start menu **or** in Add/Remove programs, try installing the latest version, then uninstalling. This might also help solve missing installer errors (like File {7c7a8eba-cb1a-4626-aa2d-8311b84092ab}.msi not found) during uninstall.

Part



Common Tasks

5 Common Tasks

5.1 Getting notified of a network problem (alerts)

5.1.1 Creating / Configuring Alerts

What is an alert?

Alerts basically monitor the conditions of a specific IP address, and then do something when those conditions exceed a specific range. The things you can do with an alert are:

- [Send an email](#)^[45]
- [Play a sound or .WAV file](#)^[44]
- [Log to a text file](#)^[47]
- [Change the tray icon / show a message](#)^[48]
- [Launch an executable](#)^[47]

For example, let's say you need to know when a destination you're monitoring stops responding. You can attach an alert within PingPlotter to that IP address so that you receive an email alert if the last 10 of 10 sample requests are lost.

Another possible alert condition to check for is if the average for the last 10 samples is above 500ms (or any other number). You can send an email alert, maybe play a .WAV file (if you're usually within hearing distance) or both. If you're trying to show your ISP there's a problem, you might log the data to a file so you have records of every time it happened over a time period.

One user had hardware problems with his cable modem. He set up an alert that [launched an executable](#)^[47] that then communicated with a device attached to his computer to reboot the cable modem when alert conditions were met. There are a whole variety of things you can do with alerts and events.

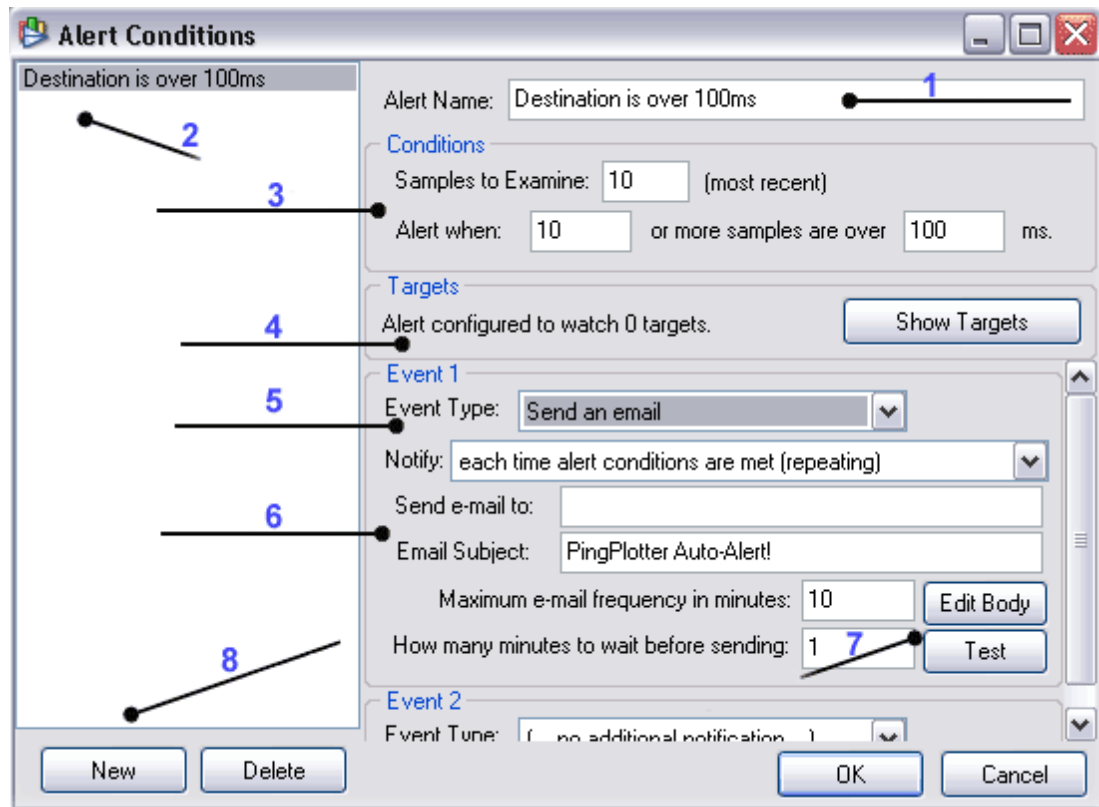
Alert setup can at first seem confusing, but it involves these four steps

1. Setting the Alert Name
2. Setting up the Condition that will trigger the alert.
3. Specifying the Event(s), or what you want PingPlotter to do when your alert(s) are triggered.
4. Tying the alert to a target that you're monitoring.

What is an event?

Events go hand in hand with alerts. Any number of events can be created with an alert – so that when some conditions occur, something happens. The "Alert" part specifies the conditions and serves as a container for the events. The "Events" specify what action(s) will happen. You can have any number of events tied to an alert, and any number of alerts tied to a host. Note: You must have an alert tied to an IP address before it will work - keep reading for more details.

Alert Dialog / Settings



Alerts are defined in the Alert Setup screen under the Edit/Alert Setup menu. Once there, you'll see an image similar to the one above. (Note: if you haven't setup an alert before, you won't see any alerts listed in the list area on the far left of the screen shot). This screen is pretty busy, but also pretty self-explanatory. From the screen capture above, you can see that we have (labeled on the image with corresponding numbers in this list):

1. The Alert Name. In this example we're using "Destination is Over 100ms".
2. The list of alerts you currently have defined. You can see from the image that we have two defined.
3. The Conditions for the alert. In other words, what has to happen for this alert to fire.
4. The Show Targets button that allows us to see how many targets/hops this alert is assigned to, as well as the IP address or DNS name for that hop or hops. If this alert is set to watch 0 targets, then the alert will not fire! We discuss this in more detail later in this topic.
5. The Notification/Trigger Events for this alert. An unlimited number of events can be tied to an alert, but most likely you won't use more than four or five. Please see the [Event Notifications/Trigger](#)

[Notifications section](#)^[44] for all of the trigger notifications available.

6. Each Event/Trigger Notification will allow you to "do something", based on the overall type of alert you're defining. For example, with a [Tray Icon Alert](#)^[48] you define the message the pop-up will contain, a [Send an Email Alert](#)^[45] needs to know where to send the email and with what frequency to send, a [Play a Sound Alert](#)^[44] needs to know what to play, etc. You enter that information here under the trigger event.
7. (These options are only for the [Send an Email](#)^[45] Event type). The Edit Body button will take you into the email Alert Email Body Editor, where you can fine tune the alert message body for optimal viewing on a cellphone, pager or other type of handheld device. By clicking this button, you can [verify that your email settings are correct](#)^[42], and what your alert email will look like.
8. To define a new alert you click on the New button. To delete an alert, highlight that alert in the list above these buttons and click on the Delete button. When you're done defining the alert click the OK button. Cancel exits this screen without saving anything.

So referring to the image above, you can see we have an alert with a Condition defined where when the last 10 samples are greater than 100ms, we want PingPlotter to email us. We've called this alert "Destination is Over 100ms ". If you were watching for a timeout, you'd enter 9999 instead of 100. Do note that 9999 is not a magic number. A lost packet is always greater than any number entered, so you can use 1000, or 20000 here and a dropped packet will exceed either of those numbers.

Setting up an alert


It's usually best to start out with an alert and event that's easy to verify - like playing a sound, or changing the tray icon. So, let's set up the alert to play a sound when packet loss goes about 40%. Packet loss is a big bad thing, and in many cases is more important than latency. We run a similar alert at Pingman HQ pretty much continuously.

- Run PingPlotter and go to the "Edit/Alert Setup" menu. If you've never set up an alert, you won't have any listed here.
- Hit the "New" button (near the bottom left).
- Enter a name for the alert (We're using "Server is Down") in the "Alert Name" section.
- In traces to examine, use "10". Set it to alert when 4 or more samples are over 9999. (Note: 9999 is not a magic number. A lost packet is always greater than any number entered, so you can use 1000,

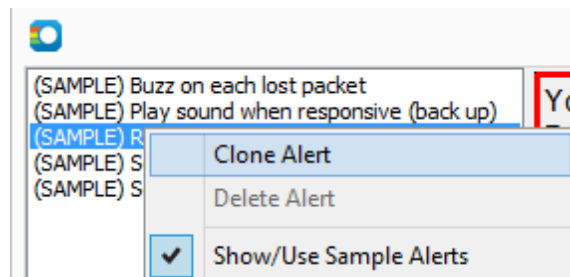
or 20000 here and a dropped packet will exceed either of those numbers). That setting will alert when 40% packet loss is achieved. The 40% comes from saying to alert if 4 of the last 10 samples were dropped. You might want to do 40 of 100, or 2 of 5 – it depends on the period of time you want it to look at. You probably *don't* want to notify if only one sample is lost, but maybe 2 in a row (which would be examining 2, and alerting if 2 or more traces over 9999). Lots of things you can do here.

Sample alerts

Starting with PingPlotter v4.10.1, "sample" alerts are included in the alert setup screen:

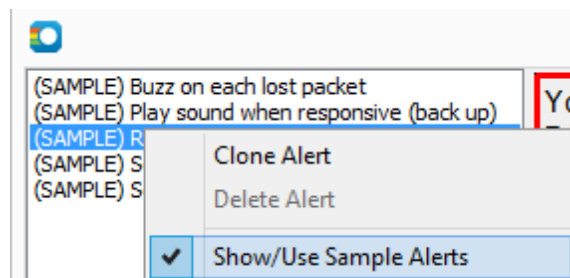


These are just a few of our favorite setups here at Pingman Tools headquarters, and are great if you're just looking for a quick alert setup. If you see an alert that you'd like you use, just right click on it, and clone it:



From there, you can tweak the configuration if you need (or leave it as is),

If you should find that these sample alerts are getting in the way of your workflow, you can always opt to hide them (and they can always be unhidden later on, if needed):



Alerting on packet loss

A common condition to want to alert on is packet loss. The fields you need to manipulate are in the Alert setup screen is the "Alert conditions" portion.

Example: Let's say you want to notify when packet loss equals or exceeds 40%.

To do this, set "Samples to Examine" to 10, and Alert when "4" or more samples are over 9999ms. A lost packet always exceeds any number you enter in the threshold area, so if you want to consider only explicitly lost packets, set this to 9999. If you want to consider any really high latency packets as well, set this to something lower (maybe 1000 or 1500).

This only examines the last 10 packets, but let's say you want to examine a higher period - and notify on a lower packet loss percentage.

Set Samples to examine to 10000 (or some other high number). Alert when "500" or more samples are over 2500ms. This will alert when you hit 5% packet loss over a period of a few hours (depending on what trace interval you use).

Picking an event

Next, select what you want to happen when the alert fires. We're going to "Play a sound" (This is in the Event 1 area). Notice that as soon as you change the "Event type" to "Play a sound", that "Event 2" will appear, with "(.. no additional notification ..)". You can have as many events as you like – and to delete an event, just change its type to "(.. no additional notification ..)".

As soon as you pick an event type, a set of options relating to that event type will come up. The "Notify" dropdown is described in more detail elsewhere. For more details on how to set up a sound event, [click here](#)⁴⁴.

For this example, use "Notify: each time alert conditions are met (repeating)" and then hit the folder icon beside the "Play Sound:" edit field. Browse to an appropriate sound file (.wav or otherwise) on your computer to fill this box.


Using this event, each time a sample is sent out and the alert conditions are met, the selected sound will play.

Tying an alert to an IP


Once you've got your alert set up, you need to tell it which IPs to "watch" for those conditions. To do this, trace to a destination just like you normally would. Once the path has listed, pick the router/destination from the list you want to monitor - and right-click on the hop number. From the popup, select "Watch this host (Alerts)...". If you want to monitor a destination that isn't responding for some reason,

just right-click on the lower time-graph for that host, and you'll have a similar popup menu.

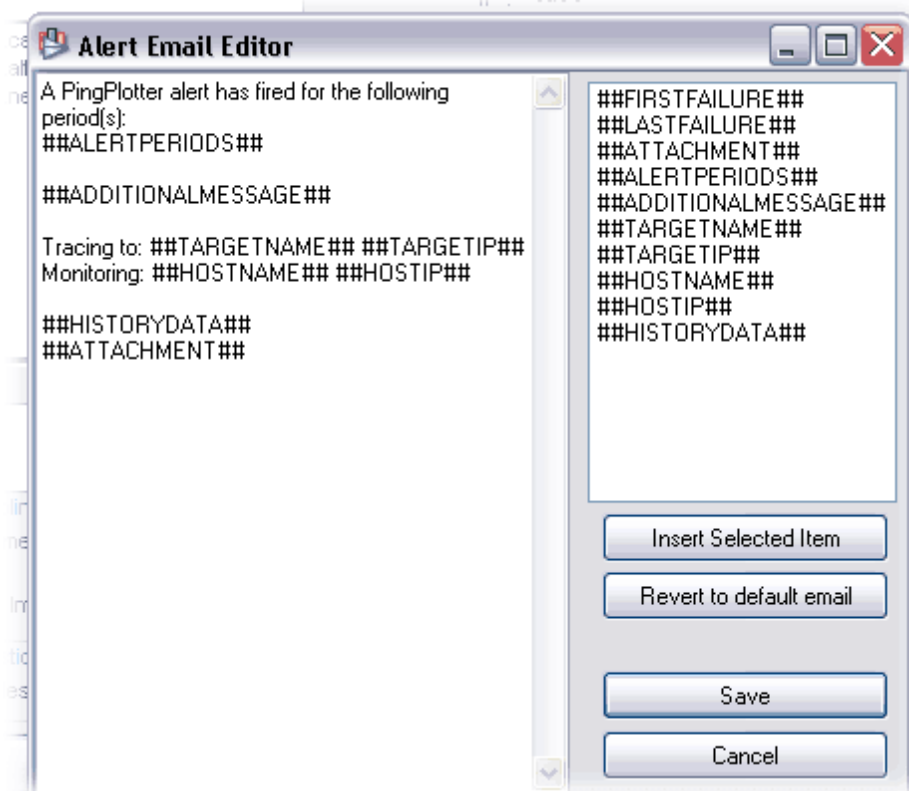
You'll get a dialog that shows the DNS name (if there is one) and the IP address you selected - and then a list of available and selected alerts. Move the alerts you want applied to this IP address into the "Selected" list - and then close the dialog. Monitoring will start immediately.



You can stop monitoring by right-clicking on the hop number, selecting "Watch this host (Alerts)...", and then removing any alerts from the selected list.



Editing the body of an email



The Alert Email Body editor allows you to fine tune the body of the alert email to whatever makes sense for your particular receiving device.

For example, if you have alert emails going to your cellphone's inbox, you may only want to have the Target Name, First instance of Failure and Last Instance of Failure in the email. If you have any more text than that, it will be too long, your cellphone will truncate it and the text is pretty much useless to you. We recognized the fact that a lot of people are using RIM devices, pagers or cellphones to receive alert email, and so now we have an editor to help you out.

You can customize the email body text to your liking in the left pane's editor, highlighting attributes on the left and clicking the Insert Selected Item button to enter that attribute into the body text to the left. This ability to highlight and insert the alert specific attributes is a great way to avoid typographical errors, and is highly recommended. When you're done with your changes, just click the Save button to save your new alert email body. If you get to the point where you want to start over again with the original email, just click the Revert to default email button and then start your changes anew. If you change your mind about editing the email body at all, just click the Cancel button to go back to the Alert Editor.

[A note about "average" response times.](#)

Average response times are a problem. The real problem with **mean** averages is when a server stops responding - what is the average of the last 10 samples if the last 10 were timeouts? Because of the problem with this we always do "when X or more samples is > Y" (this is a **median** average). You can still get your alert to work like an average - by saying "when 5 or 10 samples exceeds 300 ms" (this would be like a mean average over 300ms, but would also fire when there were lost packets).

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

5.1.2 Troubleshooting Alerts

If alerts aren't working, there are a number of things you can do to troubleshoot. Here are some suggestions - and feel free to check out our [support pages](#) for more information.

Make sure you're "Watching" a host with the alert.

By far, the most common reason that an alert isn't working is because it isn't tied to an IP address. An alert won't just start working automatically – you need to tell it which host(s) you want that alert to watch.

To attach an alert to an IP, trace to the host you're interested in. Then, right-click on the hop you want to monitor and select "Watch this host (Alerts)..." from the menu.

Note that if you're tracing to a destination that doesn't respond, and you want to watch that destination (even though it's not on the upper graph), just right-click on the time-graph on the bottom and this same menu item should be on that menu.

From here, you can move an alert from the "Available" to the "Selected" list. Any alerts on the "Selected" list will watch this host whenever this host is involved in any route (be it an intermediate host, or the final destination).

When a host is being watched by the alert system, there will be brackets around the hop number in the upper graph. If those brackets aren't showing, then that host isn't being watched. This should put a [...] around the hop that's being monitored.

Set up an alert that will fire instantly, with an event that is very evident.

If you have an alert set up – and tied to a host (see above), but it seems like the alert isn't working, then changing your alert parameters (or create a new "test" alert).

Set up "Traces to Examine:" to 10. Alert when "1" or more traces are over 1ms.

Unless your network is responding in 1ms or less, this alert will fire on the first collected sample with the alert enabled.

For an event type, use "Play a sound", or "Tray icon change/notification" as both of these events happen immediately with no wait. In addition, for the "Play a sound", use "each time alert conditions are met (repeating)", as this will continuously make sound, rather than just when conditions start / stop.

Using this sequence, you should be able to tie an alert to just about any host and have the alert conditions fire immediately. Now, add on another event type (ie: email).

You can leave multiple events tied to a single alert – that way you can continue to hear the sound while you're troubleshooting another event type.

5.1.3 Email Setup (for alerts)

The email setup dialog is used to set up emailing for alerts. If you're not using alerts, or you're not interested in having the alert system email you, then setting this up is not required.

Return Address

All outgoing emails will have a return address specified, and this is the address that is used. Please make sure you specify a valid address here since this is where all the bounce messages will come from. Some ISP SMTP servers only allow emails sent out with a "from" address of their domain as well, so if you're having problems getting the SMTP server to work, make sure you're using a valid return address.

Attach PingPlotter Savefile (.pp2) on Email Alert

When an alert goes out, data will be attached. The data that is attached encompasses all the alert period since the *last* similar alert was sent out. This is a global setting for all email alerts. Note that you'll need PingPlotter installed on a machine to read the attached data.

Include "alert" samples in text message

If this option is checked, any outgoing alert emails will include text showing the samples that failed the

alert. As each alert email goes out, the past history is marked as being sent already so that you won't get duplicate data in reoccurring emails.

SMTP Server

The SMTP server is the server that your outgoing mails will go through. This may have been given to you by your ISP or your mail administrator.

Server Port

The default port for most SMTP servers is 25. If you connect to your SMTP server via a different port (such as 587, a likely candidate if your SMTP server uses STARTTLS), then enter that port here.

Leaving this blank will use port 25.

SMTP Authentication

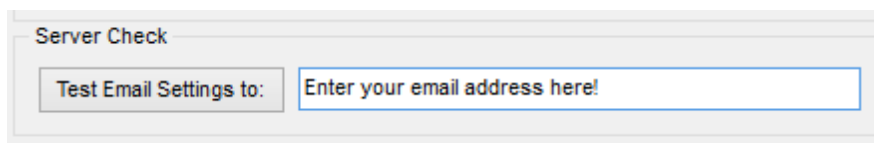
Some SMTP servers require a username and password to be able to deliver mail. If this is the case with your server, turn on the "Use SMTP Authentication" checkbox, and then enter your username and password. The password is saved in your MultiPing.ini file using a basic XOR encryption scheme – this will keep your password hidden, but this encryption method is relatively simple to "crack" if someone really wants to figure it out.

Use STARTTLS encryption if server supports it.

Most current SMTP servers prefer or demand the use of a secure channel to protect your username and password. If you use one of these servers (GMail is a great example of one), then you'll need to turn on STARTTLS. For more details on this topic, [see our knowledge base article](#).

Server Check.

Once you've got all of your email options configured, you can run a quick test to make sure everything is working as expected:



The image shows a window titled "Server Check". Inside the window, there is a button labeled "Test Email Settings to:" and a text input field with the placeholder text "Enter your email address here!".

All you'll need to do is enter an email address (preferably the one you'll be sending your email alerts to), and click on the "Test Email Settings To:" button. If everything is set up correctly, you'll get a prompt letting you know that a test message was delivered to your inbox (and you'll also get a "test"

email in your actual email inbox!).

5.1.4 Alert Events

5.1.4.1 Event Notification

Many of the events share a notification mechanism. Here is a list of the types. Note that any alert can have multiple events of the same type, so you can set up a single alert to do something at any one or all of these times.

Each time alert conditions are met (repeating)

The event will happen every time conditions are met. This means the event will happen over and over again – on each sample that causes the alert to fire. In previous versions, this was the only supported notification type.

When alert conditions start (enters alert state)

The first time alert conditions occur, the event will happen. As long as the conditions continue, though, the event won't be repeated. This is a popular use – as you find out about new conditions when they happen, but don't have to be bothered again. As soon as the alert stops happening, then as soon as it starts again, this event will happen again.

When alert conditions end (leaves alert state)

This happens when network conditions improve so that the alert is no longer firing. As soon as the conditions move from bad to good (based on your settings), then this event will happen. A use of this is to have PingPlotter email you each time conditions go bad (see above), and then when they improve again – but not to tell you anything in between.

Each time alert conditions are **not met**

This is the exact opposite of the first notification type above. As long as things are good on the network, this event will fire each time a sample is collected.

5.1.4.2 Event - Play a Sound

One of the most simple event types is to play a sound (i.e.: .wav file) of some kind.

This event can happen based on the [standard notification rules](#)^[44], and can play anything that Windows multimedia sound function wants to play. If you want to launch a sound file that this event type doesn't

support, use the "[Launch an executable](#)^[47]" option instead, as it will launch any file, including sound files.

Click the folder on the right side of the file name to browse for a file. When browsing for files, you can right-click on any sound file to play it (this is an operating system feature, and may not be supported on all operating systems).

Enter "BEEP" (no quotes) here if you just want to beep the computer speaker instead of playing through the sound card.

5.1.4.3 Event - Send Email

A very popular event type is the "Send Email" event.

Before you can create an event to send an email, you must [configure your SMTP server and return email address](#)^[42]. Note that your SMTP server must be accessible on the network to be able to send emails, so it's possible a network failure may not be able to email you. PingPlotter will continue to try to send emails once a minute until it is able to get an email out.

Emails are a bit more complicated to set up than most Event types – as it is dependant on your SMTP server, and you don't want to be overwhelmed with emails when conditions are bad, but you *do* want to know what's going on.

First off, you can fire emails based on the standard PingPlotter notification types. [See the associated documentation for more details](#)^[44].

Send e-mail to:

This can be an individual email address, or a list of addresses separated by either a , or a ; (both work equally well). Please do not set this up to be someone at your ISP unless they have agreed that they want to see this information. A huge portion of getting problems solved is playing the game right, and overwhelming people with automated emails is almost certainly going to work against you.

Email Subject:

This defaults to "PingPlotter Auto-Alert!", but can be customized with a variety of variables / text. [Click here to see this list](#)^[51].

For emails, the \$host makes a huge amount of sense (i.e.: \$host is down!), while the time/date aren't as useful because the email already contains data about this in most cases.

Maximum email frequency and minutes to wait

The next two settings control the frequency at which you'll get e-mails during alert conditions.

We recommend settings both of these values to 0 (no delay) when using a "When alert conditions start" and "When alert conditions end" event types.

When using a "Every sample" event type, though, this will result in massive overload of emails, so you'll want to limit the number of emails sent.

The first (Maximum e-mail frequency) means that you'll only get emails that often for this alert/target/event combination. Once you get an alert email, you won't get *another* one until this amount of time has passed. If conditions call for an email to be sent, it's delayed until this amount of time has passed. Beware that conditions may have changed between the time the alert is "queued" (and delayed) and the point where the alert email is sent out (because it's been delayed by this setting). This can be confusing (which is why we recommend setting this to 0 on alert conditions start and alert conditions end events).

The second (How long to wait for worse conditions) specifies how long PingPlotter will wait after its first alert condition to send an e-mail. This option allows you to wait a few minutes to find out if it was a temporary or more permanent alert condition. You may not want one immediately - because you'll want to wait a bit more info to be included - so you may want to wait 5 minutes or so before that first e-mail gets sent off.

Testing and error messages



Once you have your e-mail set up, use the "test" button to see what the message will look like (and also to make sure all the settings are working). Any errors should be displayed here.

Many of the errors that occur during testing can be attributed to incorrect email setup – [so go there and validate your settings](#)^[42].

Here are some specific knowledge base articles on possible error numbers:

[Socket error 10053](#)

[Socket error 10060](#)

[Socket error 11004](#)

Any Socket error is being generated by the SMTP server itself, not PingPlotter, so if you're getting an error number not listed here, or in [our knowledge base](#), try doing a search on your favorite search engine to see if you can find more information about the error you're getting.

5.1.4.4 Event - Launch an executable

This event type gives you all kinds of capabilities to do things when network conditions go south.

While this option is called "Launch an executable", it can actually launch any file. It can launch documents, links, .mp3s, batch files, whatever – really anything with a file association that Windows will know what to do with.

Of course, you can control when you want to launch the executable (when the problem starts? when it ends?). [See the associated documentation for more details.](#)^[44]

The filename to launch (and/or parameters for that filename) [can use variables](#)^[51] to pass an IP Address, name and date/time to the called program. Note that these are always parsed – and there is no way to cause these to be passed as literals. If you need to have one of these strings as literals passed to an executable, then you'll need to set up an intermediate link, batch file, or similar.

Note: The launching program isn't closed – it's just launched. You'll need to configure your setup to do any appropriate follow-on actions.

5.1.4.5 Event - Log to file

The "Log to file" alert writes data to a text file whenever alert conditions are met.

Log times for entire route?

This option specifies if you want to write data to the text file just for the monitored host, or the entire route. Leaving this off means that for each time alert conditions are met, one data item will be written to the file – for the monitored host. If the switch is turned on, then data for the entire route will be written.

Filename:

The filename is *required* to have the \$host variable in it ([or some variable that changes depending on which host / target is being logged](#)^[51]). If it is missing, then the file will be nonsensical if you attach this alert to more than one host. The following variables can be used as part of the filename. Note there is no way to "escape" the following sequences, so these are always parsed and can't be specified as literals in the filename.

\$dest	The target destination's name (or IP Address if the IP didn't resolve to a name). This is different than \$host because an alert might be attached to an intermediate hop, whose information is accessed via the \$host variable(s). The final destination, though, is accessed via the \$dest variable(s).
---------------	---

\$destip	The destination IP address
\$destname	The resolved DNS name for the destination
\$host	Host name (or IP address, if no name exists). If used in an alert on a non-target hop, the alerting hop's information will be used.
\$hostip	The host IP address
\$hostdnsname	The resolved DNS name for the host
\$year	Current Year
\$month	Current Month
\$day	Current day of month
\$hour	Current hour
\$minute	Current minute
\$second	Current second
\$date	Same as \$year-\$month-\$day.
\$time	The same as \$hour\$minute - note the absence of any punctuation - that's to make sure the file name is valid, if this mask is being used in a file name.

Starting with PingPlotter 3.40, directories will be created to any depth needed to satisfy your file name (prior to this version, only one level of depth was created). If you specify `c:\ppdata\%year%\%month%\%day%\%host%\%host data for %day.log`, then each directory level will be created.

5.1.4.6 Event - Tray Icon Change

This is a great event to add to most of your alerts. It's helpful to be able to see if there's an alert condition under way, and a quick glance at the tray can let you know by using the event.

The Tray Icon Change notification will do one or both of the following:

Change default icon to red during alert conditions

If you already have PingPlotter showing in the tray, this will change the existing green icon and add red to indicate that an alert has fired. If you don't have PingPlotter in the tray already, then a red icon will be added to the tray. When the alert condition(s) are over, the icon will change back to green.

Popup message in tray



This shows a "balloon" message coming out of the tray. Not all versions of Windows support this message (ie: some versions of Windows 95), in which case no balloon will show. Only one balloon can be shown at a time, so the newest balloon always wins (a new balloon message will replace an older one).

[Variable substitution](#)^[51] will be done on your message text.

One of our favorite messages is: Alert fired on \$host, \$date \$time. This is a great message because it stays up until you acknowledge it, so when you come back to a PC, you can see what alert fired, what host it happened on, and what time it *last* happened. This may be over a weekend, but the message will still be there telling you that an alert happened.

5.1.4.7 Event - Web REST Call

The Web REST call event gives you the capability to fire an alert to any service or program that has the ability to receive REST calls.

When you create a Web REST alert event for the first time, you'll notice that there's a "place-holder" example already in place:

Event 1
Event Type: Web REST call (http or https) ▼
Notify: when alert conditions start (enters alert state) ▼
REST Address: https://api.rest-dummy.test/?hostname=\$host
REST Method: GET ▼ Test
Body: example_only=true
Format: target_description=\$dest
name=value message=Alert on \$destname (\$destip)!
name2=value2
If unable to connect, retry every: 60 seconds ▼
Discard duplicate undeliverable failures:

You'll need to have the specific address for whatever service or program you're wanting to fire the alert through here. *Most* services will have some sort of documentation on where to find this address, and what needs to be included in the body (i.e.: a security token) in order for things to work.

5.1.4.8 Event - Add/Remove from Summary

This event type gives you the capability to place (or remove) targets into custom summary screens based off of their performance (which you can specify when entering your alert conditions).

The screenshot shows two event configuration panels.
Event 1:
Event Type: Add to summary screen (dropdown)
Notify: when alert conditions start (enters alert state) (dropdown)
Summary to add to: High Latency! (dropdown) with a Test button
Set focus on summary too (probably just for testing) (checkbox)
Event 2:
Event Type: Remove from summary screen (dropdown)
Notify: when alert conditions end (leaves alert state) (dropdown)
Summary to remove from: High Latency! (dropdown) with a Test button

To set up this alert event, you'll want to have a custom summary screen set up (you can create one by typing a name into the "Summary to add to" field and clicking on the "test" button").

One of our favorite setups with this alert type is to have two events (as shown in the above screenshot). Using this alert, anytime a target starts to experience high latency (we have the above alert set to look at the last 100 samples, and notify us if 50 or more are over 500ms), that target is moved into our "High Latency!" summary screen. If the target falls back under that threshold, then it's then removed from the summary screen (which is why event one is set to notify when alert conditions start, and event two when they end). This gives us a quick view into any of our targets that are currently experiencing issues.

This feature in this topic is exclusive to PingPlotter Pro. See our [product comparison](#) page for more details

5.1.4.9 Event - Add Comment

The "Add a comment in time graph" event is a great way to log any network events that may have happened while monitoring a target.

The screenshot shows the configuration for Event 1:
Event Type: Add a comment in time graph (dropdown)
Notify: each time alert conditions are met (repeating) (dropdown)
Comment to add: Alert fired on \$host. (text field) with a Test button

The setup here is fairly straight-forward: set up your alert conditions, pick a preference from the "notify" drop-down, and then enter the comment you'd like to be added to the timeline graph (hint: you can use

[variables](#) (51 here!)

5.1.5 Variable Substitution

A number of the alert parameters allow you to insert a variable which will be substituted for a value when the alert happens. Here is a list.

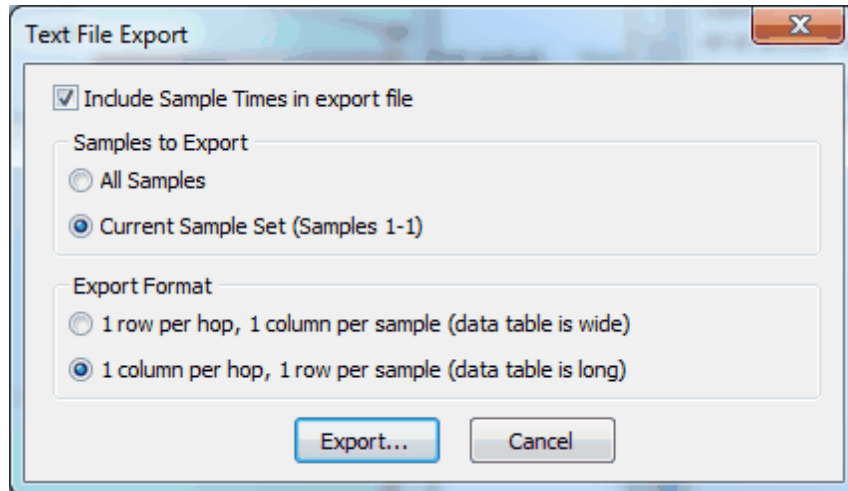
\$dest	The target destination's name (or IP Address if the IP didn't resolve to a name). This is different than \$host because an alert might be attached to an intermediate hop, whose information is accessed via the \$host variable(s). The final destination, though, is accessed via the \$dest variable(s).
\$destip	The destination IP address
\$destname	The resolved DNS name for the destination
\$host	Host name (or IP address, if no name exists). If used in an alert on a non-target hop, the alerting hop's information will be used.
\$hostip	The host IP address
\$hostdnsname	The resolved DNS name for the host
\$year	Current Year
\$month	Current Month
\$day	Current day of month
\$hour	Current hour
\$minute	Current minute
\$second	Current second
\$date	Same as \$year-\$month-\$day.
\$time	The same as \$hour\$minute - note the absence of any punctuation - that's to make sure the file name is valid, if this mask is being used in a file name.

5.2 Exporting data for further analysis

If you're looking to export data from PingPlotter, this can be done via the "File" -> "Export to text file..."

menu option. This option allows you to take all of the data currently in memory in PingPlotter and save it to a text file which can then be loaded into Excel (or a similar tool).

We're going to briefly discuss the options here. If you want more information and an example Microsoft Excel spreadsheet, check our [knowledge base](#) (in particular, [this topic](#) and [this topic](#)).



Include Sample Times in export file

If left unchecked, the times for each trace don't get saved - only the data.

Samples to Export

Select **All Samples** if you want to export your whole trace to text. Select **Current Sample Set** if all you want to save is the currently displayed sample set shown on the Time Line graph. The current sample set is the setting you'd use if you're wanting to email trace data to an ISP, etc. - though just saving a graph would be a better option.

Export Format

PingPlotter gives you two different export formats to save your data in. Both are shown below (Note: if you're using Excel, the second "1 column per hop" usually works best).

Examples:

One row per hop, one column per sample

```
,Sample Times,,9/11/2014 11:13:29 PM,9/11/2014 11:13:30 PM,
    9/11/2014 11:13:31 PM,9/11/2014 11:13:32 PM,9/11/2014 11:13:33 PM,
    9/11/2014 11:13:34 PM,9/11/2014 11:13:35 PM
1,,N/A,N/A,N/A,N/A,N/A,N/A,N/A
2,loop1.bois-dsl-gw1.bois.uswest.net,216.161.136.254,10,10,10,10,10,20,10
3,100.fa2-0.bois-agw1.bois.uswest.net,207.108.229.29,30,20,10,10,10,10,10
```



```

4, ----- ,207.108.224.247,20,10,20,10,20,10
( ... middle data snipped for brevity ... )
16,192.ATM7-0.GW5.SJC1.ALTER.NET,152.63.54.21,60,51,50,50,50,50
17,digexoc12-gw.customer.alter.net,157.130.214.154,50,50,50,50,50,50
18,gsr-01-p2-0-a00a02.af.sjc5.digex.com,164.109.130.26,50,50,50,51,50,50,50
19,lc1.com,164.109.154.154,50,50,50,51,50,50,50

```

One column per hop, one row per sample

Host Information

```

1,,
2,loop1.bois-dsl-gw1.bois.uswest.net,216.161.136.254
3,100.fa2-0.bois-agw1.bois.uswest.net,207.108.229.29
4, ----- ,207.108.224.247
( ... middle hosts snipped for brevity ... )
16,192.ATM7-0.GW5.SJC1.ALTER.NET,152.63.54.21
17,digexoc12-gw.customer.alter.net,157.130.214.154
18,gsr-01-p2-0-a00a02.af.sjc5.digex.com,164.109.130.26
19,lc1.com,164.109.154.154

```

Sample Information

```

"9/11/2014 11:13:29 PM",N/A,10,30,20,30 .. snipped .. 60,50,50,50
"9/11/2014 11:13:30 PM",N/A,10,20,10,30 .. snipped .. 51,50,50,50
"9/11/2014 11:13:31 PM",N/A,10,10,20,30 .. snipped .. 50,50,50,50
"9/11/2014 11:13:32 PM",N/A,10,10,10,30 .. snipped .. 50,50,51,51
"9/11/2014 11:13:33 PM",N/A,10,10,20,30 .. snipped .. 50,50,50,50
"9/11/2014 11:13:34 PM",N/A,20,10,20,30 .. snipped .. 50,50,50,50
"9/11/2014 11:13:35 PM",N/A,10,10,10,30 .. snipped .. 50,50,50,50

```

****Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details****

5.3 Long term monitoring and auto-saving data

If you run PingPlotter 24 hours a day, 7 days a week, you're going to want to save your data (in case of power failure or other event), and you're also going to want to limit the memory footprint for PingPlotter. We talk about some best practices and recommendations here, although your needs may be different. For example, if you're using PingPlotter and monitoring a lot of targets, you'll probably want to keep fewer samples in memory for each target, otherwise the footprint of PingPlotter will be pretty big.

Setting up your memory footprint

PingPlotter defaults to keep 250,000 samples in memory. If you regularly do long term monitoring, though, you may want to understand this number so you can change it to fit your needs. In particular, if you keep too many samples in memory, you may run out of system memory at some point.

- Determine how often you want to sample. 2.5 seconds (the default value in PingPlotter) gives a good amount of accuracy without too much data. Some people do use 10 seconds (although, anything

much longer and you might miss problems).

- We find that having 4 to 5 days of data in memory at a time works well. There are 86,400 seconds in a day, 432,000 seconds in 5 days. Divide this 432,000 by your trace interval. For 2.5 seconds, this gives us 172800 samples in 5 days.
- In PingPlotter, go to "Edit" -> "Options" -> "Auto-save" section. Enter your calculated number in "Maximum samples to hold in memory". 172800 samples takes up roughly 10 to 15 megs of RAM in memory, which puts the PingPlotter memory footprint around 40 megs total (it keeps multiple copies of the data in RAM at some points, and general overhead). This is workable for just about any workstation, unless you're tracing to more than a handful of targets.
- For PingPlotter Pro: If you're using multiple configurations, make sure you review the proper settings for *each* named configuration.

Setting up PingPlotter to save data

With 4 to 5 days of data in memory, each save of data will have all of this - which puts each save file around 1 to 3 megabytes. Having one file per day gives you easy access to a day's data, along with the previous 4 days for good analysis. We suggest saving every 30 to 60 minutes, with a file name like this:
c:\ppdata\\$dest\\$dest \$date.pp2

- In PingPlotter, go to "Edit" -> "Options" -> "Auto-save" section.
- Turn on "Auto-save data"
- Set "Save Interval" for "30 minutes"
- Set file name to "c:\ppdata\\$dest\\$dest \$date.pp2". If you're currently tracing to a target, floating over the file name field will show you a hint of the file that would be actually saved.

Note: It's very important that you specify an absolute path for your save file name! If you're running as a Windows service and you get save files in your Win32\System directory (or some other unexpected directory), that's because you didn't set an absolute path here.

We set up "30 minutes" for a save interval. The file name controls how often we create a new save file. If the file already exists, then it will be overwritten to include the new data. You can include \$hour in the save file name to get a new file every hour, but we don't recommend this for save data, since you'll get a new save file every hour and you may run out of hard drive space.

This will give you a new file each day with 5 days of data in it. Each day's file will be missing the last few minutes of the day as the 30 minute save interval may hit at 23:35 or 23:59, but that data will always be

stored in the *next* day's file.

Feel free to tweak these settings however you want. This is just a discussion of some possible starting points.

Autosaving with a lot of targets

If you're using PingPlotter Pro and monitoring a lot of targets, you'll need to be cognizant of the impact of auto-saving a lot of targets. Each time the auto-save timer fires, PingPlotter completely overwrites each of your auto-save files. With a lot of data in memory and a lot of targets, this can take a few seconds. As you approach the memory limit of a machine, this can often mean memory swaps occur as well, to pull in the old samples and write them out to a file. On some machines, this can take a minute - during which time tracing can stop and the GUI can become unresponsive.

The best way to manage this is to limit the number of samples in memory - if you can keep just a couple of days in memory for each target you'll still have the history and you'll get better performance (and have more available memory on that machine).

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

Part



Understanding Output

6 Understanding Output

6.1 Discovering a route between you and a target

Let's run through a basic scenario where we discover the route between your computer and an interesting destination (maybe one you're having a problem with).

Normally, when troubleshooting a problem, you want to run PingPlotter against the server where you're experiencing problems. Maybe that's a web server (in which case, you'll want to enter that web server's address); maybe it's a game server (in which case you'll want to enter that game server's address). If you aren't experiencing problems with your network connection or something you access with it currently, no worries! Just try and think of something you access regularly with your Internet connection (such as www.PingPlotter.com. or www.Google.com) to use in the exercise below.

We're assuming here that you've downloaded and installed PingPlotter. If you haven't, please see the [Downloading and Installing PingPlotter](#) ¹⁶⁷ section for instructions on how to do so.

So load up PingPlotter, and let's get started!

1. Enter the IP Address (i.e. 129.41.62.29) or the DNS Name (i.e. www.PingPlotter.com) of a destination you may be having problems with into PingPlotter's Address to Trace: input box. Note: we just want the name of the destination. You would not enter <http://www.problemserver.com/index.html> here. What we want is between the "http://" and the "/index.html."

For now let's leave all the other settings you see on the screen as they are.

2. After you've typed in the address, either press the "[Start](#)" ¹⁶⁷ button, or press the "enter" key on your keyboard. The trace then starts, and you'll see the upper Trace Graph populate with the route information to the target you entered. The Timeline Graph for that target will be displayed also below the Trace Graph.

One thing that is kind of an 'ah ha!' moment for a lot of first time PingPlotter users is seeing that you really do have that many devices your network traffic passes through to get to web sites, servers, etc. If you click on a web page link, that 'click' is passed on by all those hops to that final web server/page, that web server executes that click, passes the information back to you through all those hops and you see it on your browser.

It's important to cover the concept of a "[Sample Set](#)," because we're going to be mentioning it a lot. The sample set is defined by the "[Focus Time](#)" ¹⁶⁷ value (which is right under the "Address to Trace" and "Trace Interval" fields). The "PL%" and "Avg" columns in your trace graph are all computed off of this number. If this value is set to 10, PingPlotter uses the last ten samples it's done and bases everything off

that number. If it's set to 20, PingPlotter uses the last twenty traces it's done, etc. As we go through what you're actually seeing on the graphs, just remember that the "Sample Set" is based off the value in the "Focus Time" value, and that number represents the number of samples - starting with the most recent and working backwards.

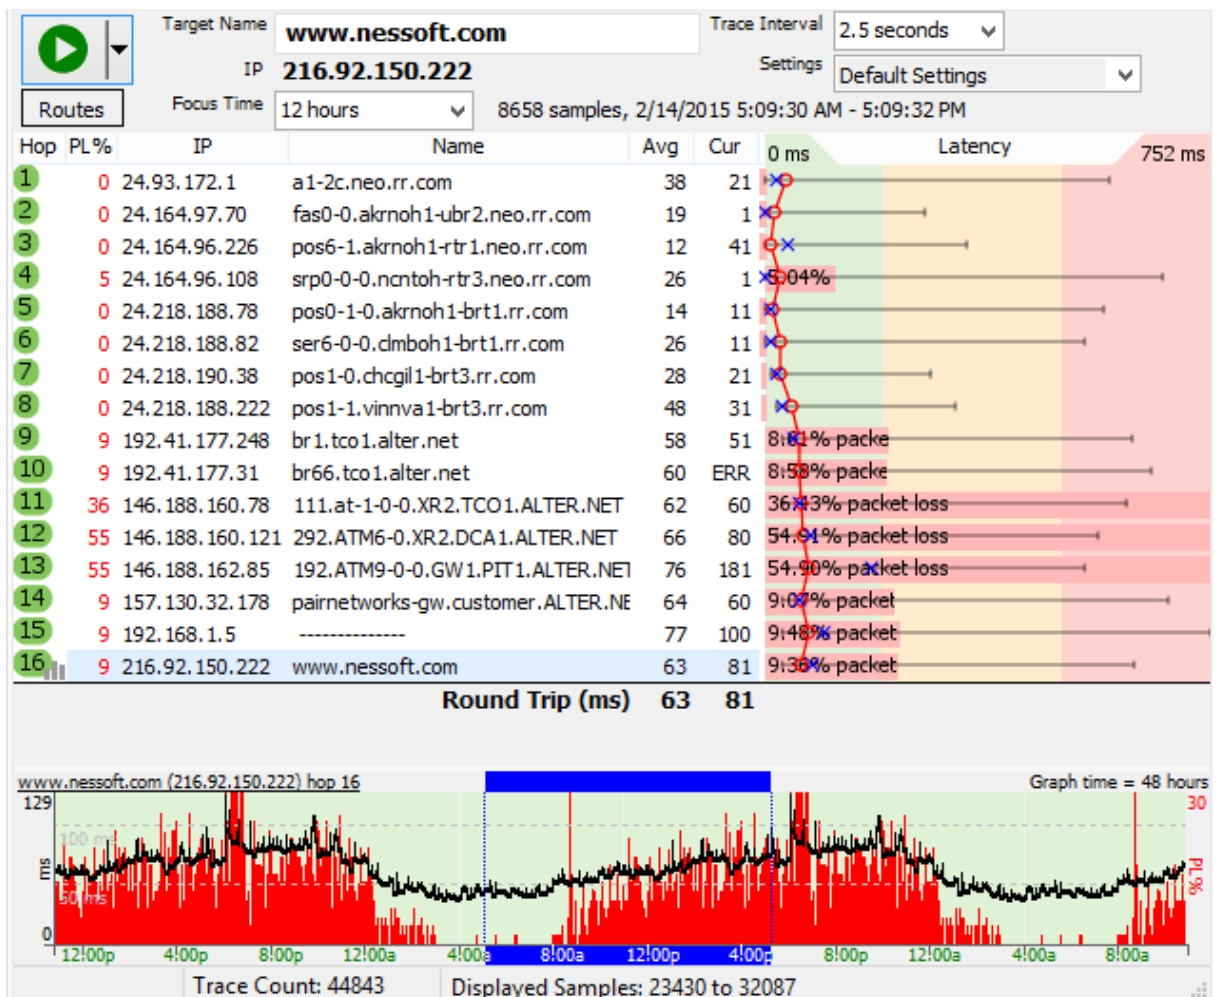
Some things to consider before we move on:

- We cover what's "normal" for latency and packet loss in a [knowledge base](http://www.pingman.com/kb/42) entry at <http://www.pingman.com/kb/42>
- If you get a "Destination Address Unreachable" message at the last hop in the trace graph, this means something between your computer and the final destination isn't receiving and/or returning packets. We cover this in detail in the PingPlotter section of [our knowledge base](http://www.pingman.com/kb/8) at <http://www.pingman.com/kb/8>. If some hops are responding, you might try using a different target address (i.e.: try [www.PingPlotter.com](http://www.pingman.com) instead of the address you entered the first time).
- If the final destination is working (i.e.: the Round Trip row is showing), but some of the earlier hops are not, then don't despair! This could be normal. See knowledge base articles <http://www.pingman.com/kb/24> and <http://www.pingman.com/kb/29> for possible causes of this behavior.
- If you're entering an IP address and would like to "label" that address, or give it a "friendly name" to make it easier to find in history later, you can do that by entering the IP Address, then a space, and then the label. For instance, if you have a core router that you'd like to see displayed in PingPlotter as "Core Router" (minus the quotes), you'd enter it as "192.168.0.1 Core router" (omitting the quotes and substituting the 192.168.0.1 for the IP Address for the actual IP Address for your router). For more details on this, see <http://www.pingman.com/kb/23>.

6.2 Interpreting Results - A Quick Example

So let's get into some specific examples of how to interpret the results from PingPlotter.

For the first example, you're getting intermittent packet loss to www.nessoft.com. What can we determine from the graph below?



First off, the final destination (hop 16) shows 9% packet loss. There's a problem someplace in the route, but we need to determine where....

Hop 4 shows 5% packet loss. Hop 5 doesn't show packet loss, though, so you know that the problem in hop 16 isn't because of hop 4. Hop 4 is likely just a router using a different CPU path for TTL=0 packets than it does for routing data through.

Hop 9, however, shows 9% packet loss, and this packet loss is carried on through to the final destination. This is a **huge** indication of where the problem lies.

Now, all we know from this is that the problem happens after hop 8. We don't know if it actually happens because of CPU overloading in hop 9, a router problem in hop 9 (or even on the exit side of hop 8), or if it's the connection between hop 8 and 9. A little bit more troubleshooting is needed for this.

Digging deeper, we can see (from the domain names) that hop 8 is in the rr.com domain, while hop 9 is in the alter.net domain. Also, the IP addresses show decidedly different ranges. This is a strong clue that

it's actually the connection between hop 8 and 9 that's causing the problem. It's likely that there's not enough bandwidth between those two locations.


6.3 Finding the source of the problem

We learned how to do a basic PingPlotter trace in the [previous exercise](#)^[57]. What we want to do now is take a pre-prepared save file that contains about 2 ½ days worth of fictional data and do more hands-on work with PingPlotter.

First off, you need to get the data file downloaded for this exercise. The file you need to download is at www.pingplotter.com/gsg/www.nessoft.com.pp2. Save this file to your desktop so it will be easy to find (of course you can select another location - just remember what you specified).

Load up the save file in PingPlotter by going to "File" -> "Load Sample Set," and then browse to your desktop (or wherever you saved the nessoft.pp2 file) and select the www.nessoft.com.pp2 file, then click the "Open" button. Ta da! Two and a half days worth of data for us to play with.


Let's go over some concepts concerning the Timeline Graph.



1. Let's start out with this concept – red on a graph is bad. All those red lines you see on the Timeline Graph indicates Packet Loss for that particular time period. Remember, that means that PingPlotter didn't get an answer back. There are times where you could have a flaky router, or even a server that is de-prioritizing ICMP packets. If you'll remember from the [How PingPlotter Works](#)^[12] section, this is how PingPlotter gets its information. You could have a perfectly fine connection through that router, but PingPlotter will show it at 100% packet loss. For the most part though, when you see red on the Timeline Graph it means PingPlotter wasn't able to get to that server.
2. The black line on the Timeline Graph is average latency for the time period you're looking at. When you zoom in on the graph throughout the next two steps you'll be able to see the individual pixel-wide points on the graph.
3. The default Timeline Graph scale is 10 minutes - but what if we want to see more? Right-Click with your mouse button on the Timeline Graph, and you'll see that you can show anywhere from 60 seconds on up to 48 hours worth of data within the Timeline Graph. Go ahead and select 6 hours. Notice that you're now looking at six hours worth of trace data.

Feel free to select different time periods and see how the contents of the Timeline Graph reflect what you select. What we want to illustrate here is that via the Timeline Graph's right-click mouse menu, you can easily zoom out to look at a trend for a day, and then zoom in to a spot that looks interesting – down to 60 seconds worth of data.

4. The "focus area" on the Timeline Graph (which is brought up by double clicking anywhere on a timeline graph) shows you the current sample set that you're viewing.



Remember that value is set in the "Focus Time" value. When you first load up this save file, PingPlotter is starting at the last trace done before the data was saved, going back 10 traces, and then calculating the values for the Trace Graph based on that value. Note also that PingPlotter shows you the date and time information for the sample set that you're looking at in the area up above the Trace Graph.

You can change the focus of the Timeline graph, and subsequently the data in the Trace Graph, by double-clicking on any area of the Timeline Graph you want to look at. In case you're wondering, you can have an active trace going while you move around the Timeline Graph. You do not need to save the data, stop the trace or do anything else for that matter. Just double-click on the area you want to see.

Now change the Focus Time to 150 samples. See how the values in the Trace Graph change to reflect those additional 150 samples/traces? Change the Focus Time to ALL, or 0 (0=ALL in this context). Whoa! If you look at the Sample Set Time above the Trace Graph you'll now see that you're looking at a lot of data. In fact, you're looking at all 44,843 traces, or ALL of the traces that were in the file we loaded. Your PL% and Avg columns in the Trace Graph now reflect all those traces. Now you can change it back to something smaller (100 is generally a good value to use).

5. If you're not there already, right-click on the Timeline Graph and set the time you want to see to 6 hours. To move back through the data you click-drag, or left-click, and while the left button on your mouse is held down you move the graph to the right. To move forward in time, you click-drag to the left. Are things starting to "click," as far as how you can move back and forth through time on the Timeline Graph? If you have a lot of samples to move through, do yourself a favor and set the graph scale accordingly (again, by right-clicking and choosing a time) before you do. For instance, if you've zoomed in to where you're looking at 5 minutes worth of data and then need to go back twelve hours – set your graph scale to something like 3 or 6 hours before you start.

Tip: You can also use the keyboard to navigate backwards/forwards, as well as quickly get to the

beginning or end of Timeline Graphs. The shortcut keys used to do this are covered in the [Interface Graphs](#) ^[17] portion of this manual.

Hopefully by now you've been clicking all over the Timeline Graph, and seeing that the Timeline Graph does indeed change when you select a new sample set to look at. So how do you get back to the current time/date? Bring up the Timeline Graph's right-click menu and select "Reset Focus to Current." If you don't see that listed in your right-click menu....well, then, you're already current.

At this point if you're thinking, "Hmmm, this Timeline Graph stuff is cool, but I want to see a Timeline Graph for Hop 1 too!", then today is your lucky day. You can do this by double-clicking on that (or any) hop - or by right-clicking and selecting "Show this Timeline Graph." Double-clicking again will hide that same graph. If you have a lot of Timeline Graphs open, clicking once on the hop for that Timeline Graph will cause that Timeline Graph to flash briefly so you can pick it out of the list.

Checkpoint

So let's summarize what we've done, and what we know at this point about the Timeline Graph.

- You can change the scale by right-clicking and selecting one of the values listed.
- The focus area on the Timeline Graph shows you the current sample set, and you can change it easily to reflect the time period you want included in the Trace Graph.
- You can move around within the Timeline Graph by click-dragging or using the shortcut keys.
- By changing the graph scale, click-dragging then changing the scale back, you can zoom in and out of different time periods you want to analyze.
- PingPlotter tries to keep your "focus period" in view on the timeline graph, and you can use this if you're zooming in on a period. Focus the period you're interested in (double click on it), and then zoom in.
- If you want to get up to the very last sample set in a save file, or the last sample done if PingPlotter is actively tracing a target site for you, "Reset Focus to Current" on the right-click menu for the Timeline Graph will get you there.
- If you want to turn on or off visibility for a Timeline Graph for a hop, you can do so by double-clicking that hop, or by right-clicking on that hop and selecting "Show this Timeline Graph".
- To quickly find a Timeline Graph for a particular hop, if you have a lot of Timeline Graphs showing for instance, click once on that hop and the associated Timeline Graph will flash.

6.4 Finding the source of the problem - part 2

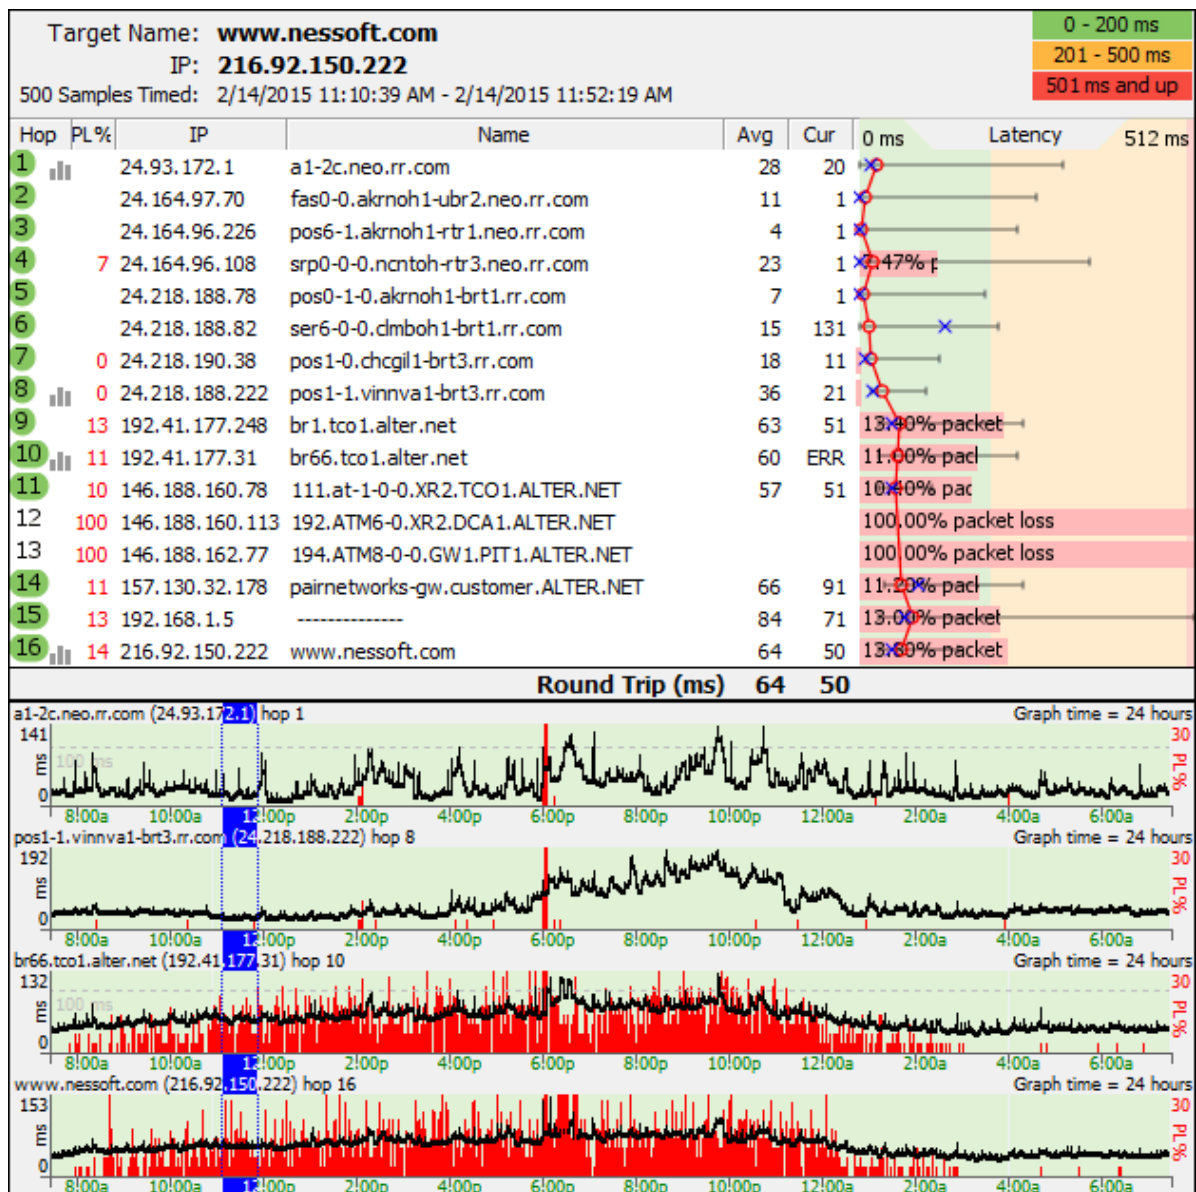
Let's apply and build on what we've learned and dig deeper.

The black line on the Timeline Graph shows you the latency over time. As you zoomed in on the graph earlier you saw that really what you're seeing is pixel-wide entries that represent the round trip time, or latency for the host/device who's Timeline Graph you're looking at.

Red, on the other hand, represents packet loss (PingPlotter didn't get anything back). We have numerous articles that go over different reasons why you could see packet loss in the Knowledge Base at www.pingman.com/kb. Our goal here is to help you understand at a basic level how PingPlotter can help you figure out why that loss is happening.

Packet loss, of course, isn't the only thing that can cause poor performance for your particular application. The other big factor is latency. For online games, for instance, latency is a killer. Let's review our sample data, and look at a trouble spot so we can learn more how to interpret what PingPlotter gives us that no other tool can.

Let's look at the time period between 11am and 12pm on February 14, 2015. Let's take the approach that we were browsing the web at that time and were running pcAnywhere into our computer at work. The pcAnywhere session goes to pot.



1. Set "Samples to Include" to a number between 100 and 500 (we're using 500 here because it's easy to see the focus period on the 24 hour graph).
2. Change the Timeline Graph scale to 24 hours (remember: right-click then pick 24 hours), and then scroll all the way to the beginning of the graph
3. Take a look around at that time between 11:00 am and 12:00 pm. Double-click a time period somewhere around the 12:00 pm time. You'll notice the focus lines appear on the time graph, and the upper graph will show packet loss.
4. Wow, look at that packet loss! 13% at hop 9, 14% at hop 16. Notice how the packet loss is added at hop 9, and then all the downstream hops also show high packet loss. This is a strong (compelling, in

this case) indicator that hop 9 (or the link between hop 8 and hop 9) is the reason we're seeing packet loss.

5. Notice the packet loss trending over time. This indicates some kind of time-based load problem. Also, the fact that the packet loss starts at the junction point between rr.com and alter.net indicates a possible problem at the connection between these two providers. It's possible that rr.com doesn't have enough bandwidth to service needs.

6. Let's turn on a couple more timeline graphs. Double-click on hop 9 and hop 8. Notice the difference in packet loss. Notice, too, that the latency is still relatively high, even at hop 8, and it's also showing time-based problems. Double-click on hop 1, and notice that that *to* also shows a time-based latency problem. This is a separate problem from what we're seeing at the hop 8 to hop 9 junction point! Looks like we should contact rr.com and find out if they can help us solve this problem!

7. Try zooming in on a problem-period by changing the timeline graph scale to something lower. When you see an interesting period, double-click on it to "focus" it, and then change the timeline graph scale to zoom in or out. This gives you a close-up view of the data with more detail.

Now, it should be noted that spikes on your graph may not be really high latency. A spike in the graph is high latency for that period displayed on the graph. The Timeline Graph auto-scales itself, so to see what that high spike is you need to look to the left-most part of the Timeline Graph. In this case the value we're looking for is 87ms.

This was a really quick example of moving around within PingPlotter and digging out information about not only what's going on with a connection, but really focusing in on a problem. Let's move on now and talk about collecting this data over time and preparing a good case for your ISP (or someone able to solve this problem!).

6.5 Interpreting Results - Longterm Monitoring

PingPlotter allows you to use the timeline graphs to zoom in on any particular time, so even if you weren't there (or didn't save an image) when something was happening, you can still recover that exact image later. You shouldn't ever have to be sitting in front of your computer when an outage happens, or you experience other problems on your network, to get the data you need from PingPlotter.

The Scenario:

You're having outages (or situations you want to communicate to your ISP) randomly throughout the day (let's say twice a day). The problem is that you can't be there every time an outage happens so you can [save a graph image](#)¹²³.

For this example, you're keeping 24 hours of data in memory or even more. We suggest that you normally use 2.5 second trace intervals and keep 200,000 samples in memory (this is almost a week's worth of data). You can change the number of samples to keep in memory in the [Auto-Save](#)^[119] section under the [Edit/Options menu selection](#)^[123].


Using PingPlotter's [timeline graph](#)^[17], you can see over the past X time period (see below) to identify a time period where there was problems. Problems will demonstrate themselves as packet loss (red), or high latency.

Now you want to show the route, and the packet loss/latency in the upper graph for that time period. Since that time has already passed, you need to change the focus of the upper graph to that time in the past.

First off, you need to make sure your "Focus Time" focuses in on just the period in question, so let's change that to 100 (it's important to **not** have it set to 0/all, but to have it be a number smaller than the number of samples in memory to be able to focus the upper graph). Right click on the timeline graph and pick a reasonable period of time to set the viewable time period. For instance, you may want to set it to six hours so you're not scrolling forever. You can then "click and drag" the graph to the left to go back to the time period you want to focus on and drag it to the right to go forward in time.

Note: [You can have custom timeline intervals show up in your right-click menu by adding a setting in the PingPlotter.ini file](#)^[100].

Double-click on the "problem period" in the lower graph. You'll see a focus rectangle appear on the lower time graph, and the upper graph will change to represent the data you have "focused" in the lower graph. Once you've done this, you might want to change the scale of the lower graph to show more detail. Right-click on the lower time graph again and change the scale to an hour (or maybe even 30 or 10 minutes depending on how long the outage was). The focus rectangle should still be visible. You can fine tune the data being displayed in the upper graph by double-clicking on the lower graph again.



Using these techniques, you should be able to zoom in on exactly the right data to best illustrate the problems you're seeing. You can look at the data after the problem occurs and get the perfect picture and not have to sit there watching PingPlotter all day and night.

You can auto-save that data by for instance having the [auto-save](#)^[119] function in PingPlotter create new files every day, and then load up a prior day to do the same thing you did above for a particular time. This gives you the capability to have pretty close to 100% coverage of your network performance and be able to zoom in on any particular outage, period of slow response, etc.

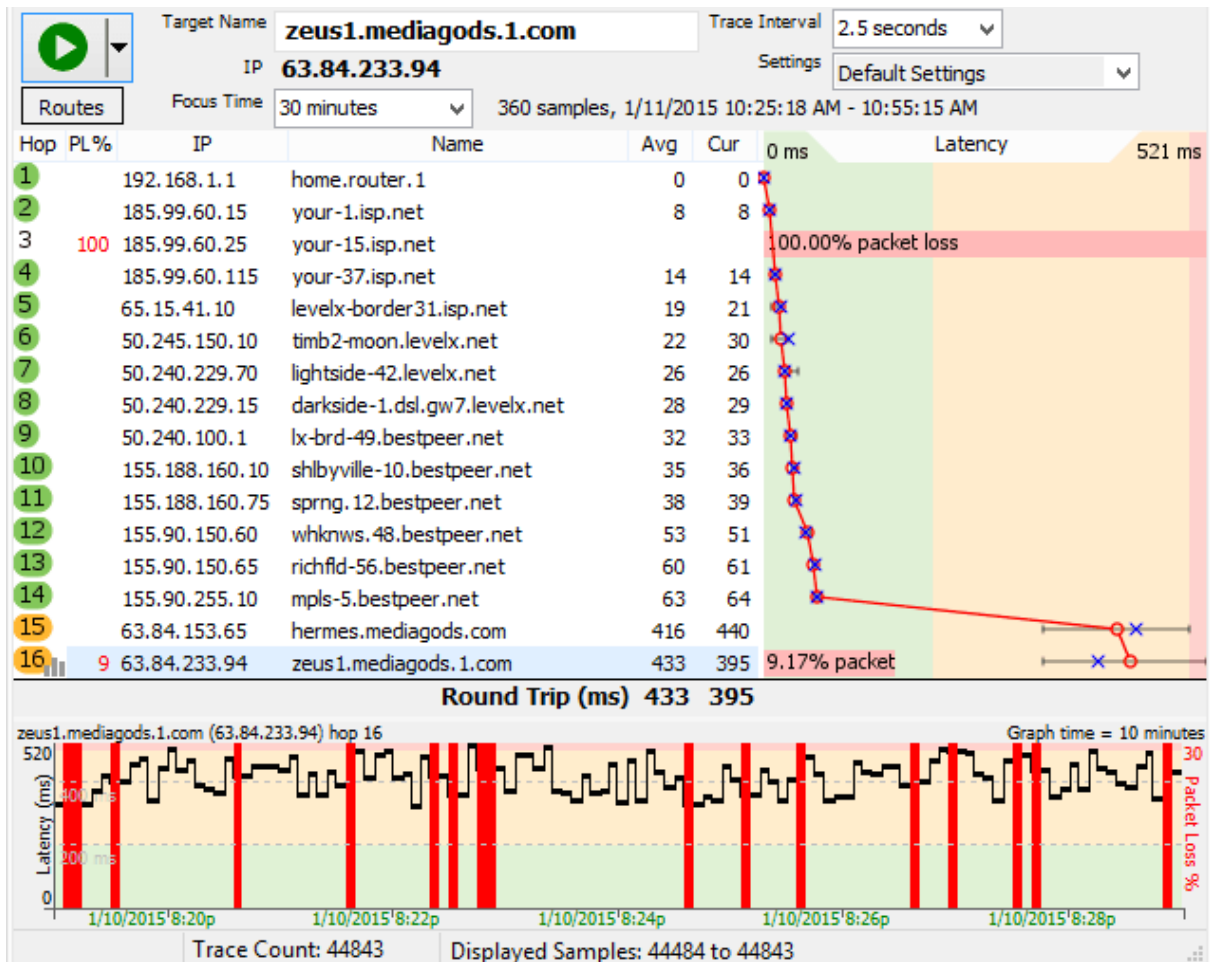
The options in the [alert setup](#)^[33] do allow you to have the .PP2 file (trace data) emailed, and then you can use these same capabilities to zoom in on that data.

6.6 Interpreting Results - Gamers

For this example, let's assume you're an online gamer - specifically, a Quake III player (though the following is representative of any online game really - MMORPG, RTS, racing sims, fighting games, etc.). You've got two servers that are running the same maps you like to play, so the only issue you have is which one out of the two is going to give you a better connection. We realize some folks aren't going to be so patient as to use the method below to decide which server they're going to play on... but bear with us here. We're learning! The same topics we go over in this section, as far as graph interpretation, are also applicable if you were trying to figure out why a connection to a specific server you were just playing on is so cruddy.

The first thing you need is the IP addresses or DNS names for the two servers. You'll launch PingPlotter, enter the IP address of the first server into the [Address to Trace](#)^[16] box and click the Trace button. 2.5 seconds is a good value for the [trace interval](#)^[16], and the [# of times to trace](#)^[16] should be unlimited (we're gonna watch it for awhile). Then, start a trace to your second server's IP address using the same settings you used for the first server.

You then go get yourself a Diet Pepsi, do some stretches or whatever in preparation for a night of gaming. When you get back to your computer you'll have two graphs that look similar to the ones below. Let's analyze both and see which server you want to play on.



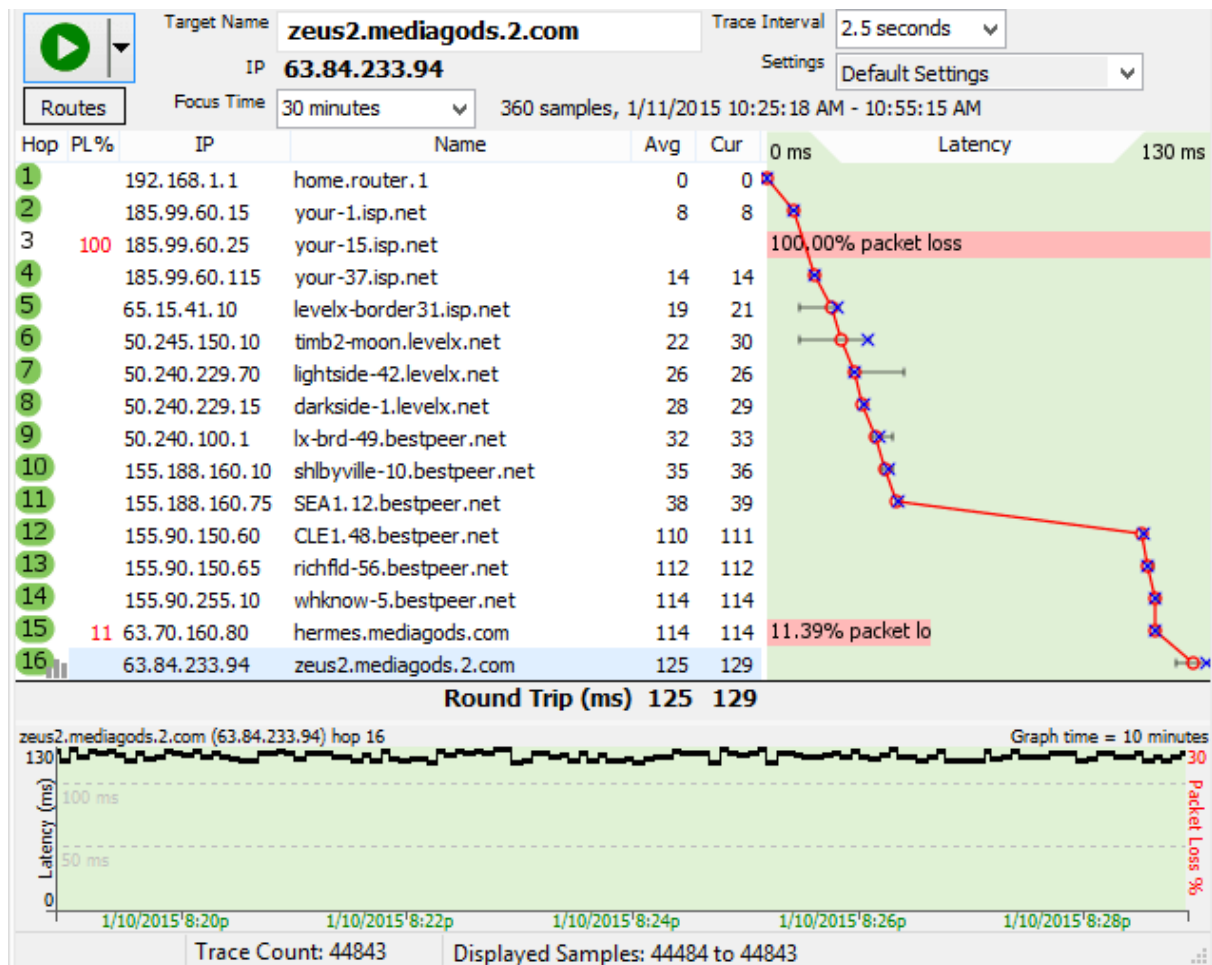
Hmm. This doesn't look too bad until you get to hop 15 and start looking at the history graph. Let's take the [time line graph](#) first.

Red is bad. Every time you see red on the history graph you had a [timeout](#), or in other words there's dropped packets. Packet loss is the bane for most online games. When you're running across a big open area and then all the sudden *blip,* you're on the other side (and most likely dead), that was more than likely caused by timeouts, or packets you didn't get to (or back from) the server. Besides the red lines on the history graph, you can also see your packet loss in the PL% column and, if you look at hop 16, the horizontal red line contains your packet loss value.

Digging a bit deeper, you can see that you're running under 100ms all the way down until you hit hop 15. Notice that you move off of LevelX's backbone into BestPeer between hops 8 and 9. No problem there, there's plenty of bandwidth between those two providers since the time doesn't really go up. From the DNS on hop 8, we can see that hop is a gateway (thus the "dsl-gw7" part of the name) to some DSL customers. Where we start running into problems is when we get off of BestPeer and hit the Mediagods domain that's hanging off the DSL link. All the sudden your latency goes up to 400+ms at hop 15. That

DSL connection is busy. Once you make it to the server at hop 16, not only is your latency still way up there, but you're getting 9 to 10% packet loss. That server is a busy bee also it seems. So busy that he's not keeping up. Combined with the bandwidth saturation we're getting on the DSL line itself, it's best to try later. We don't want to play here.

Now let's look at our second server.



Now this is more like it. Really, anything under 150 ms is a great connection. We don't even make it over 40ms until hop 12. Sweet.

Let's look at that connection between hops 11 and 12. Notice from the DNS names that you actually go from Seattle across Global Crossing's backbone to Cleveland. When you factor in speed of light latency, you can account for about 40-50ms of your latency to hop 17 with that hop across the backbone between hops 11 and 12. So you've got a 120-130ms ping to hop 16. That's pretty good. If you didn't have that fat pipe installed (and were instead running a modem) you'd probably be running at about 220-230ms for your latency.

"What about that 11% packet loss at hop 15?", you ask. Judging from the numbers for that hop and hop 16, what we're most likely dealing with here is a router that probably has a low priority for ICMP packets. A lot of network admins will set a router up to drop ping/ICMP packets first if it starts getting busy - [have a look here for more details](#)

So which server are you going to play? Obviously it's the second server above.

Other considerations

Your graph results can be affected by a number of things that are out of PingPlotter's control.

1. In the analysis of the second graph above, we mentioned that hop 15 is more than likely just a problem with that router not giving us back good information. Many routers put a low priority on [ICMP traffic](#)^[12]. Others don't even echo back ICMP requests (this will show up as a blank entry for that hop). Obviously, PingPlotter has no control over these situations.
2. PingPlotter can't track the route that your traffic takes on the return trip from the server back to you. If your inbound traceroute traffic is taking a different route back to you than the outbound traffic to the server, this is called an **asymmetrical route**. By definition traceroute doesn't take these types of routes into account and, unfortunately, PingPlotter isn't going to be able to tell you about problems with the return route in these cases. One clue that this is happening is that you'll have a great trace up until the last one or two hops on your trace. In other words, you don't have an easily identifiable problem at hop X further up the route that is mucking up the rest of the route downstream.

One thing you can do is save your trace data to a text file and post it up on a support message board for the particular game that you're playing. Even better, save off a graph and post it instead. Many sawy game server admins will actually do a traceroute (or even better a PingPlotter trace... *smile*) back to you and be able to tell you if there's problems with a route back to you when asymmetrical routes are involved.

There are some sites that can do traceroutes back to you if you want to investigate on your own. They can be found [here](#).

In closing, we can't emphasize this enough: **latency is the bane of online gaming**. Much more so than bandwidth limitations. The good thing is that PingPlotter can tell you this latency, and provide you with ammo in the way of trace and graph data when you're beating up on your ISP.

6.7 Interpreting Results - ISP Problems

For this example, we're assuming the role of a user that's having problems with a broadband connection.

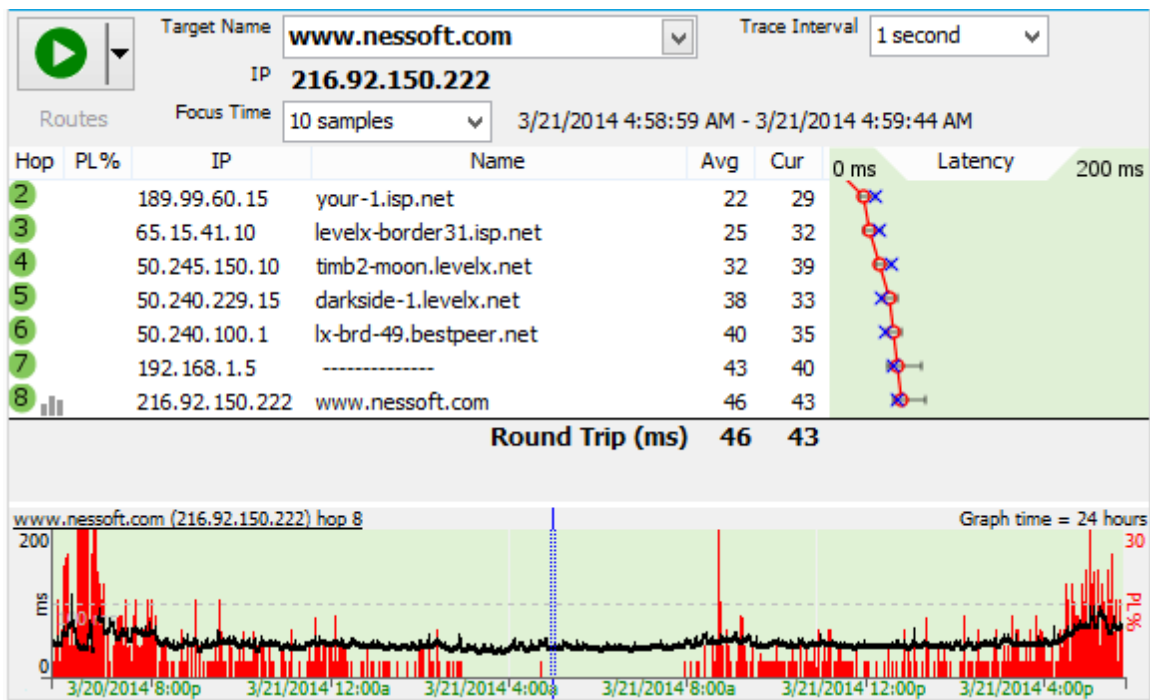
What we'll be taking a look at is a few days worth of trace data. One thing to keep in mind is that if you're doing long term monitoring and want to look at more than the largest default time span on the time-interval graph (48 hours), [you can add custom time intervals](#)^[100] in the pingplotter.ini file located in PingPlotter's installation directory.

Before continuing, if you're not familiar with how the graphs work in PingPlotter please make sure you've read the [introduction to graphs](#)^[17] earlier in this tutorial.

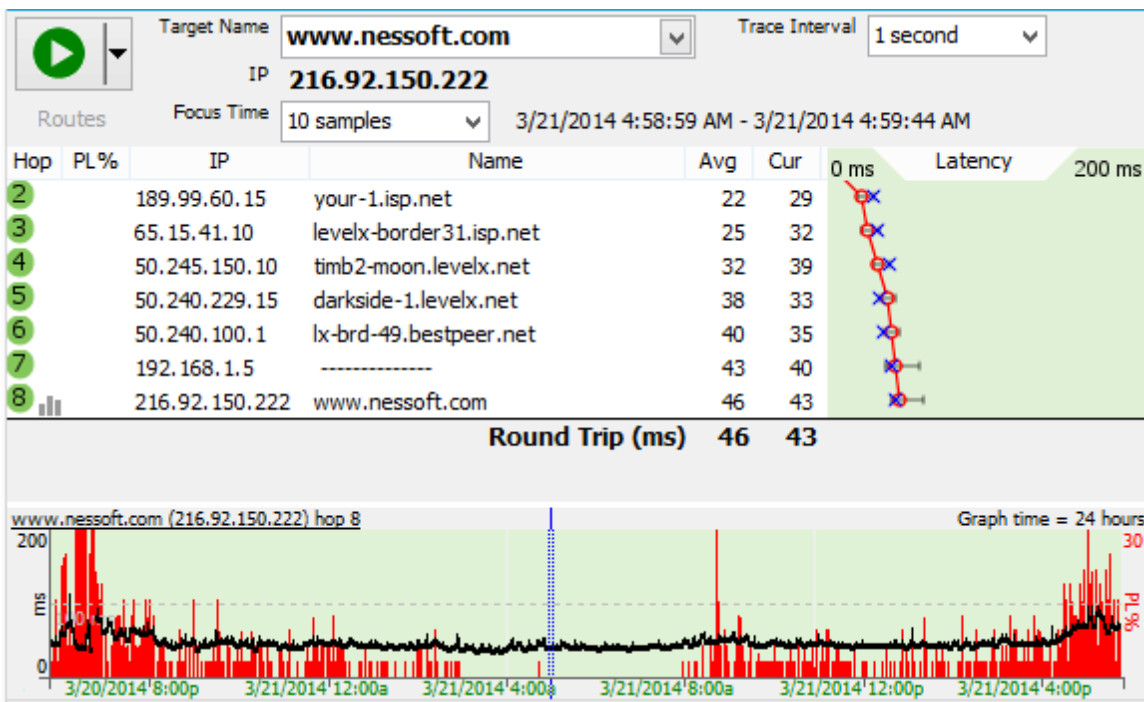
One common mistake we see folks make is that they'll trace to their ISP's border router. This is a bad thing. If you're tracing to the border router and your route changes (i.e. they take that router down for maintenance or you get load balanced onto another router) you really have no idea what happened. If you want to keep your traces local to your ISP, trace to an address that isn't going to change on you like you're ISP's mail server. This is actually a good thing to do if you're having mail problems and it's your ISP's mail server going down. Otherwise just pick a destination that you know has a reasonably good chance of always being up. This is a better choice since routes within your ISP can change, and PingPlotter keeps track of those route changes. The cool thing is that you're doing a traceroute here, not a ping, so even if that destination host goes down you can drill down on the timeline graph and see if it's your connection, or if it's just the destination being down (as in all hops but the destination don't show timeouts).

Another problem with tracing to your ISP's border router is that your ISP will not *respect* the data that you collect this way. No application targets a border, so they have no reason to trust this data. For best results, you want to pick a target that is one you use and are having problems with.


Note: For clarity, all the graphs below show us [ignoring Hop 1](#)^[124] which you can do from the View Menu. All the graphs were saved with the File/Save Image command within PingPlotter then converted to .gif for this tutorial.



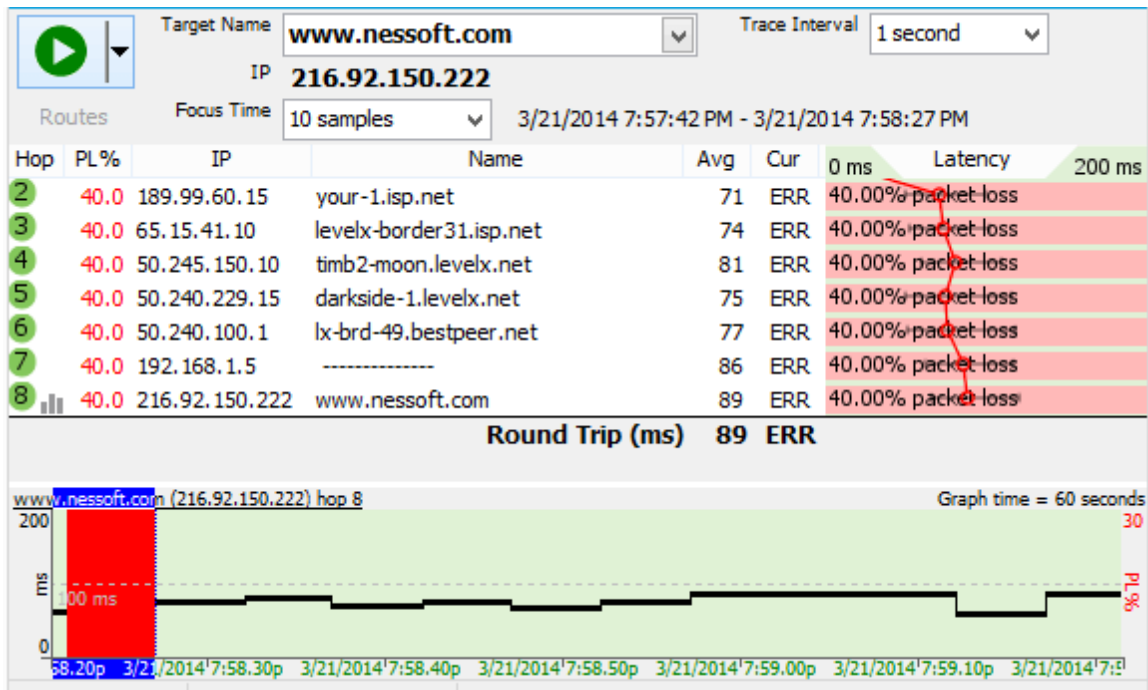
This first graph shows what the traceroute should look like with no load on the connection, i.e. no downloads, streaming audio, on-line game playing, etc. "What about all that red on the history graph? I thought red was bad?", you ask. Actually **red is bad**, however before we saved out this graph we **double-clicked on the timeline graph to drill down, or zoom-in**, and are looking at the data for 3/21/14 at 4:58 a.m. If you look at the top of the graph you see the "Focus Time" is set to "10 Samples," and the range is between 3/21/14 4:58:59 AM - 3/21/14 4:59:44 AM. So basically the above graph's trace for that *particular time* looks good. However when you look at the timeline graph, you can start to see the tale of woe. What we have here is a really flaky broadband connection. So how do we prove it? Read on.



Just sending your ISP a graph with red lines isn't very convincing. However, when you start zooming in on those sections with timeouts, and send graphs of them as well like this second saved graph, it's pretty obvious the connection's hosed when you can't see out to Hop 2. This is the same time interval as the first graph, just showing a different period in time for trace data.



For our third graph we've got data for early on the second day of our trace. Lots of red, and when we focus on the 9:01 AM time period the connection's still poor. We're unable to see out most of the time. It's hard to argue with the graph. Also keep in mind that what we're showing here is that the whole timeline graph isn't solid red. This isn't an issue where you accidentally kicked the plug on your router.



For our fourth graph, you can see from the timeline graph's times that we've adjusted the time-interval so we're only looking at 60 seconds worth of data. The trace graph is showing the section of the time-interval graph that we double-clicked on which is 7:57:42 PM to 7:58:27 PM. So what's up with the 40% packet loss showing up on the trace graph? Notice that we didn't have timeouts for about a full 50 seconds (out of 60) on the timeline graph. Out of the ten samples we're looking at, 40% of them were timeouts. This is important. When we're looking at the trace data we're looking at those 10 samples we selected and the numbers for those samples, not the whole range of data shown on the timeline graph! This is not a graph you want to send to technical support. All it's going to do is confuse them.

So in summary, PingPlotter allows you to show your ISP where the problem's are. In the these examples, we were essentially showing the whole link going down. However, we could've just as easily seen if the ISP's connection to the Internet was down at Hop 4, because we were tracing to a destination not on our ISP's local network. If there was indeed a problem at Hop 4, we would've had good trace data at Hops 2 and 3, timeouts at Hop 4 and possibly no trace data past Hop 4. If the router at Hop 3 was being flaky, and for instance you saw a lot of packet loss, it's easy to save an image showing just that so you can email it. When sending graphs to your ISP, we've found it's best to send one graph showing data for an extended time period, and then drilling down on the timeouts and sending graphs that truly show them what's going on. PingPlotter allows you to save in .png or .bmp format. We recommend .png because they're smaller.

6.8 Interpreting Results - Bad Hardware

Scenario: External customer has problems using your network resources.

A customer (not inside your network) has problems losing connecting to services inside your network. In this scenario, you are acting as a service provider for some network service. This might be provided via HTTP, or possibly through something like Citrix or Windows Terminal services.

Your customer (possibly employees of your company, or maybe subscribers of your service if you're an ASP-based service) is complaining of frequent disconnects, and possibly slow performance sometimes. How do you troubleshoot this kind of a problem? Where is the problem; at the entry point to your network, in the customer network, or possibly in one of the providers in between?

One way to pinpoint the problem is to have your customer run PingPlotter against your service. They can easily download PingPlotter, capture data, and then either analyze that data themselves, or email the data back for you to analyze. This information can be used to pinpoint which hop (or router) in the chain is adding latency or losing packets that might be causing problems back to your ASP.

Analysis:

Here's an example extracted from a real-world customer situation. We'll walk through some of the symptoms, collected data, then analysis and how we came to the proper conclusion.

Our ASP (we'll call the ASP "CitServeCo" - a totally fictitious company) is accessed via Citrix, which is relatively sensitive to high latency and packet loss. The customer (who uses financial applications served by CitServeCo) was frustrated by disconnects during the day, which lasted anywhere from a few seconds to a minute or two, interrupting their ability to do business.

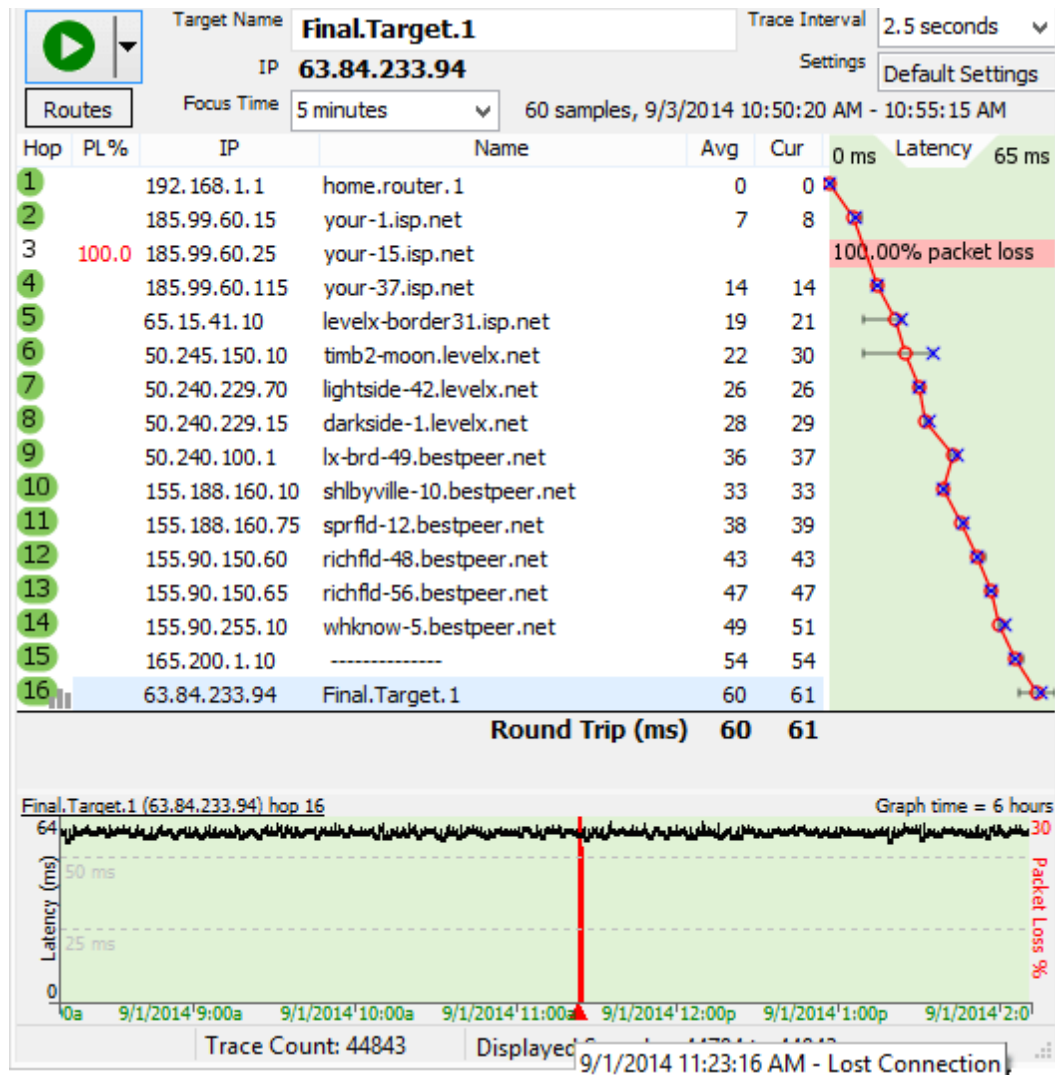
The customer will almost certainly blame CitServeCo initially for a problem like this – but we need a way to determine where the problem is being caused, and help the customer solve the problem. Any reliable ASP will get numerous complaints like this – that they *know* are the fault of something beyond their control (like the customer's cable connection, or similar). Most ASPs will attest to the fact that customer connections are a prime source of network problems, but we can't just tell the customer "It's your problem."

When the customer contacted CitServeCo, they suggested that they download PingPlotter, install it onto the workstation they normally work with, and configure it to monitor the Citrix server inside CitServeCo's network.

If a disconnect or slowdown occurs in Citrix, we suggest to the customer that the right-click on the lower “time-graph” in PingPlotter and [create a comment](#)^[27] at the point that the problem occurred. This correlates real world symptoms with the collected PingPlotter data (a crucial part of the troubleshooting process).

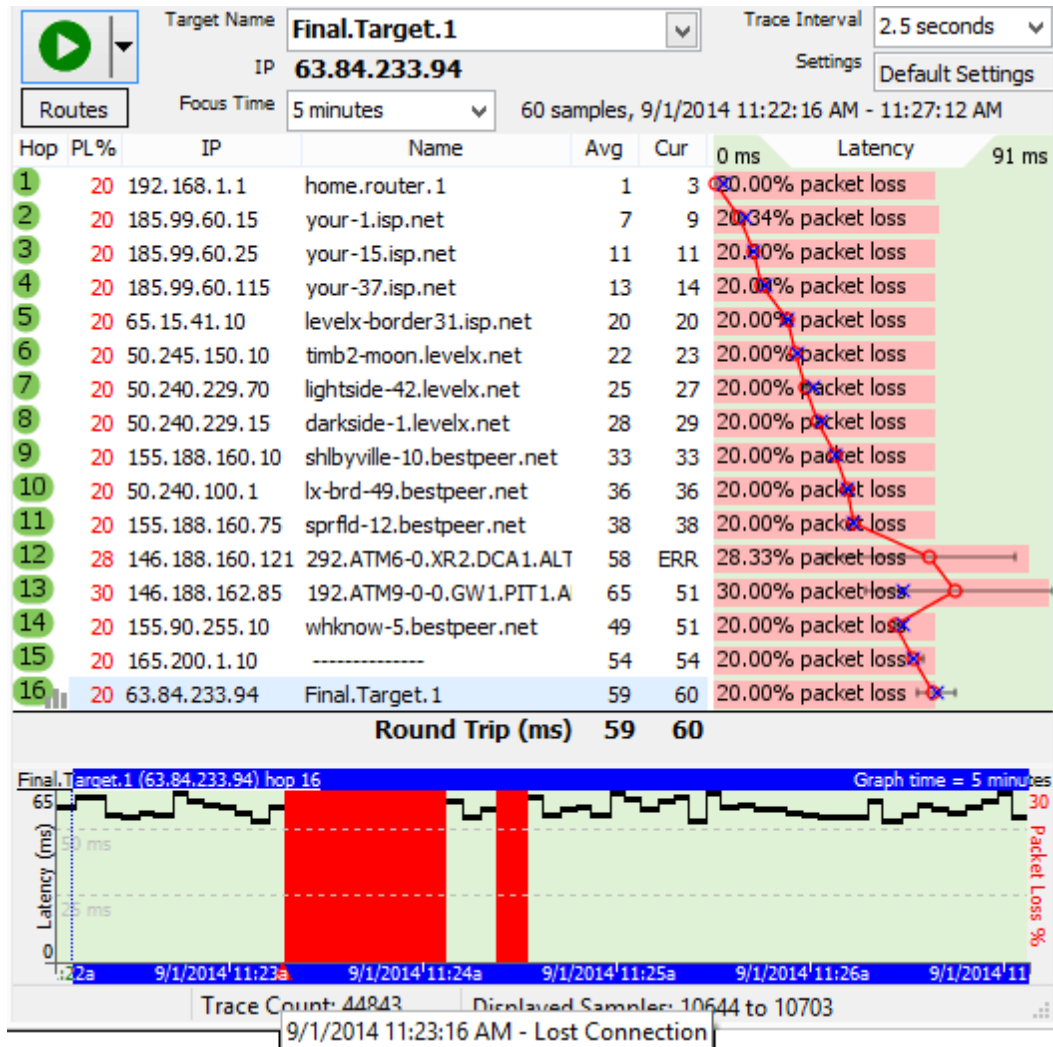
A day of collecting data resulted in only a single disconnect, and the customer dutifully recorded that for us. They then sent us the saved data for analysis. Because PingPlotter data is already compressed, there is no need to .zip up the file, making it easier to send data.

Here is the result of that day.



Notice the packet loss in the timeline graph at 11:23 am. This corresponds with a note from the customer in the data file that shows the disconnect happens. We decide we want to take a closer look. To do this, we set the “Focus Time” to 5 minutes, double-clicked on the red point in the time-graph, and

then reset the time-graph period to 5 minutes instead of 6 hours.



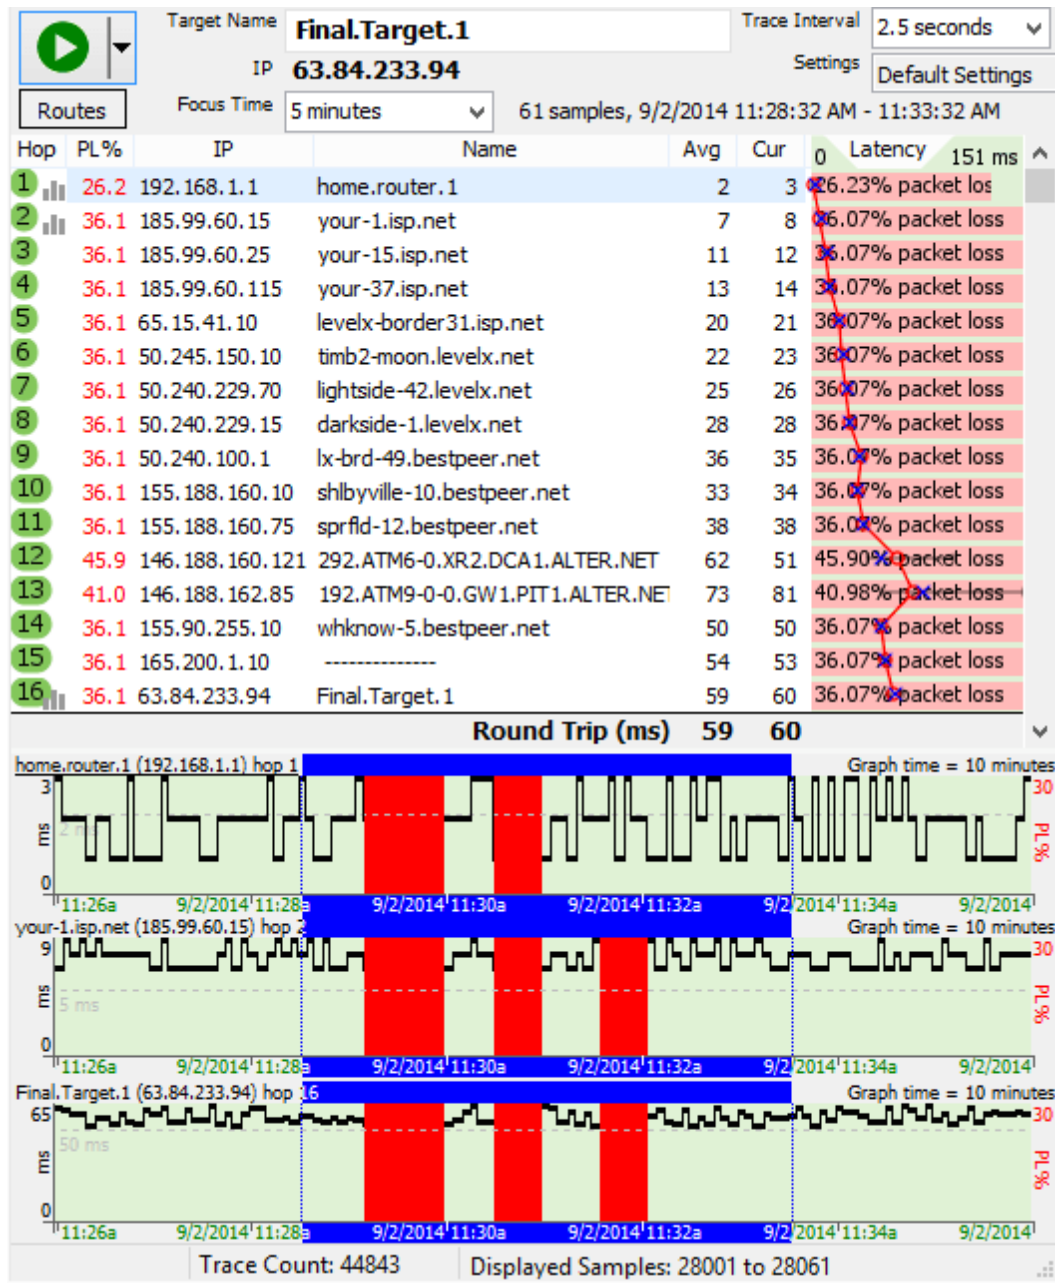
Wow. A big outage. If no packets were getting through, we're definitely going to see problems. Where did that loss occur? If we look at the upper graph, we see that we have consistent packet loss across *all* the routers (double-clicking on the lower graph focused the upper graph on the period we were interested in). The packet loss was similar, but let's have a look at the actual lost packets. To do this, double-click on a router in the upper graph to show a time graph for that router.

It looks like hop 1 lost just as many packets as hop 16. Since every piece of data needs to go through hop 1 to get to any of the other hops, a blockage there will look just like this. It looks like we've found a likely culprit – the router at hop 1, or possibly anything between the computer collecting data and hop 1. This might be as simple as a network cable, or it might be a significant amount of network equipment. We don't know until we check with the customer to find out what's here.

We asked the customer what kind of network hardware they have in place. While they were collecting this information for us, we had them continue to monitor their connection.

It turns out that all the computers at this location are hooked up to a "SOHO" router. This SOHO router is, in turn, connected to a cable modem provided by their ISP. From their ISP, the customer is uncertain as to the network configuration. We see some of this in the PingPlotter graph – a list of routers that are participating in sending data.

The customer was uncertain as to what the individual pieces of hardware reported in as, so we had them continue to monitor, but asked that they cycle the power for the devices if they had a disconnection problem. Several days later, they had an opportunity to do this, and captured the experience in PingPlotter.



In this instance, the customer had a disconnect, and then rebooted their SOHO and their cable modem. Notice how hop 1 stays working through the reboot of their cable modem, but not through the reboot of the SOHO. So this must be the SOHO. Also, notice during the disconnect, hop 1 was non-responsive. This indicates the problem was similar to having the SOHO powered off.

Maybe there was a bad power supply on the SOHO, or maybe a bad network cable someplace? The customer replaced the network cables, and the problem persisted. They then replaced the SOHO router

and the problem was solved.

It's not every case where the problem is so "in our control" as it was in this one. Sometimes, the problem requires talking an ISP into replacing some piece of their hardware, or replacing an ISP altogether. PingPlotter can be just as effective locating the outages, packet loss, or latency on ISP equipment, though – it's can sometimes just be harder to fix those problems.

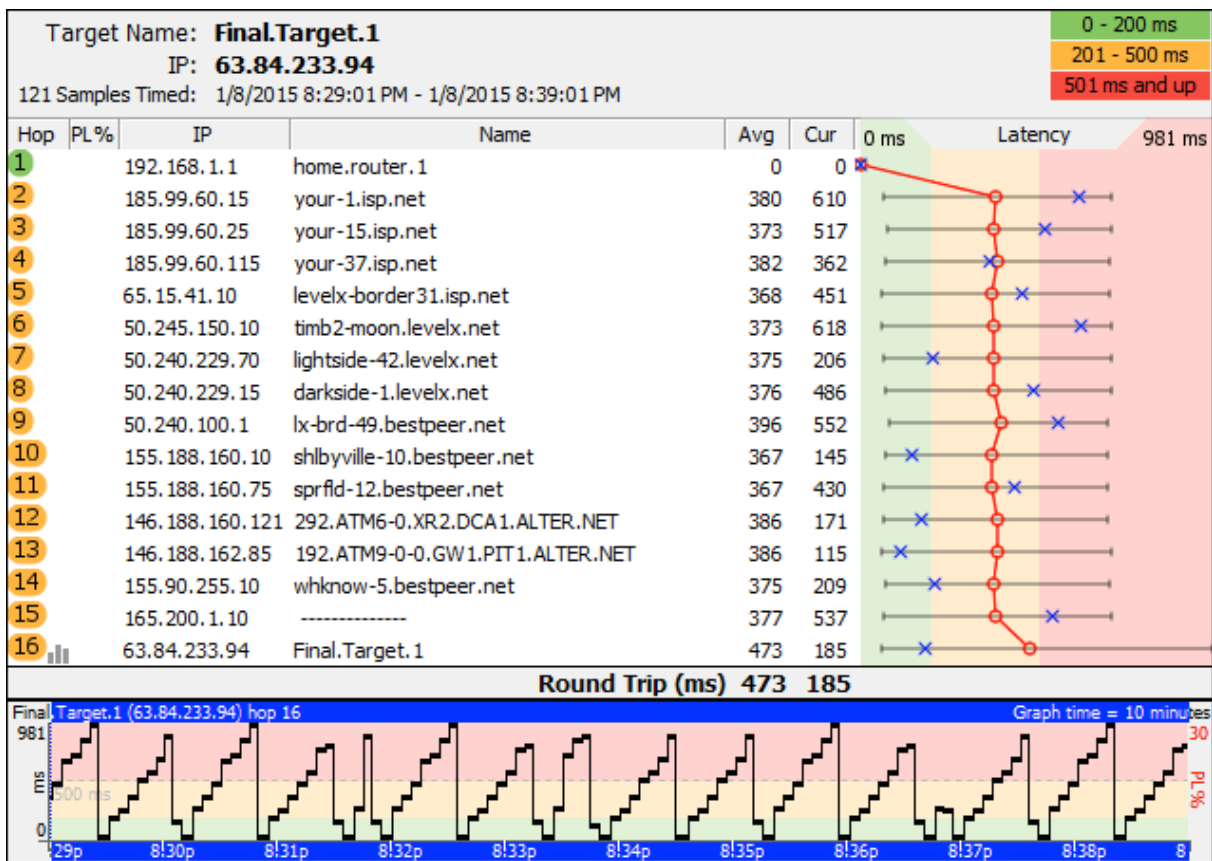
6.9 Interpreting Results - Bandwidth Saturation

Every network connection has a limit, and in this example we'll talk about how to recognize bandwidth limits on a local DSL / Cable modem.

The scenario here is a 30Mbps downstream, 10Mbps upstream cable modem running a 2.5 second trace interval to our target. The computer running PingPlotter was connected to the cable modem via a wireless network card, which loses packets occasionally.

We're downloading a 2 gig file, and during this period the bandwidth of the cable modem should be completely saturated. This means that anyone else using this network connection is going to notice significant latency and possible packet loss. Any user might decide they want to troubleshoot this situation, and might run PingPlotter to do so.

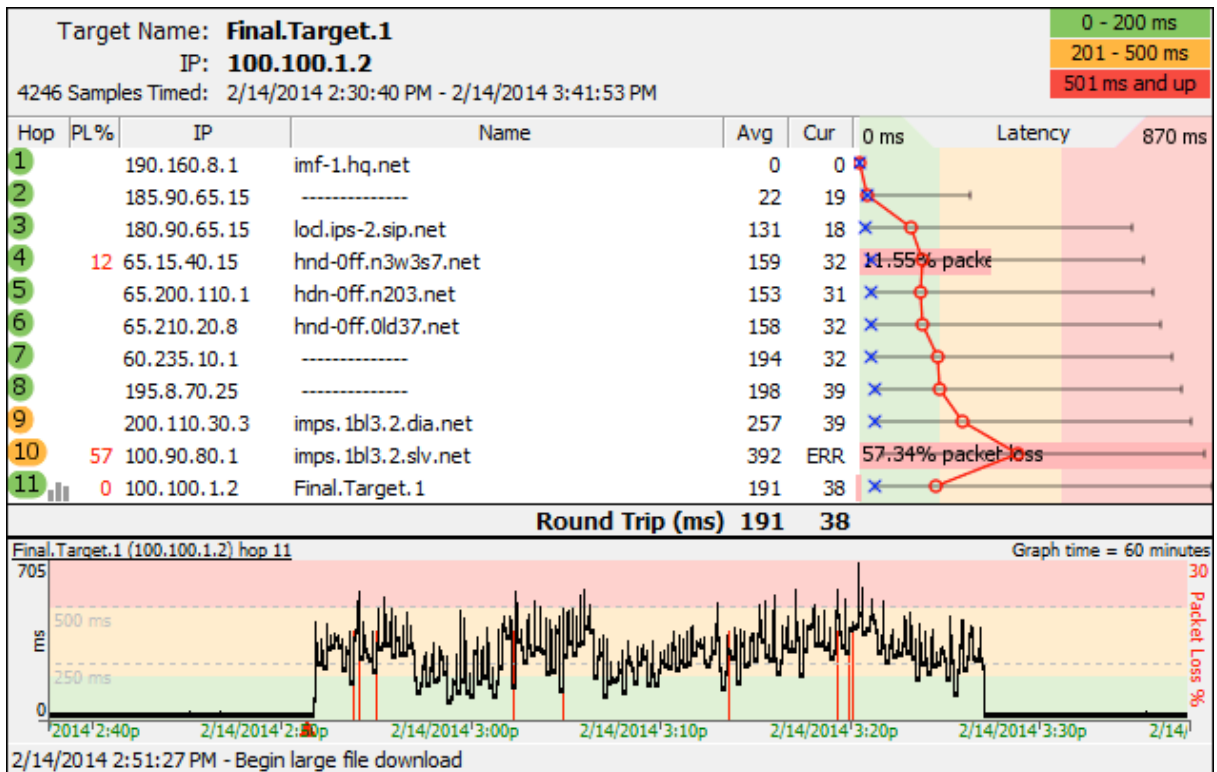
A quick 10 minute trace might look something like this:



First, notice the big latency jump between hop 1 and hop 2 - the DSL modem. This same latency jump is translated into all downstream hops, which indicates that hop 2 (or the link between hop 1 and hop 2) is significantly impacting network latency. Since we know what's going on here, it's pretty easy to recognize. In a lot of cases though, you might not know what's happening and you'll see latency like this.

Now, the "Sawtooth" pattern in the timeline graph is a classic bandwidth saturation pattern. Anytime you see a pattern like this, think "bandwidth saturation". Keep in mind that bandwidth saturation is totally normal, and happens on just about every network. Consumer-grade DSL and cable modems have the highest chance of seeing a pattern like this, but it can happen on any connection where the bandwidth at a network point is being totally saturated.

Here is another example of bandwidth saturation (with more time shown on the graph):



Here you see an 40 minute period where a large download was happening. Outside the download, the latency is pretty much rock solid. The packet loss stays relatively constant across the entire period, indicating that the packet loss isn't being caused by bandwidth saturation. We know that it's actually caused by a wireless network that has about 1% normal packet loss, and the download rate doesn't affect it at all.

There are several solutions to the latency problem:

- Procure (buy) more bandwidth
- Don't transfer as much data
- Implement QoS

This same situation happens in any network. Your ISP runs into these same exact options with the connection between their network and the outside Internet. They can limit what you download, or they can buy more bandwidth. On a bigger pipe the jump in latency isn't going to be as pronounced, and the bandwidth will be saturated by a number of different loads. Often on a primarily consumer-use network (DSL / Cable), the increase in latency will be gradual as more people use the bandwidth, and then it will gradually drop off as people stop using it. The rise in latency often happening at 6-9pm, and the drop off happening at 11pm - 1am. It's often accompanied by packet loss as well.

On your own network, the bandwidth saturation could be happening for a number of reasons: another user downloading something; temporarily restricted pipe size (problem with your Internet connection), p2p applications in use, or possibly even a virus/worm type application that is using bandwidth sending out new instances of itself.

In summary: Saturated bandwidth can be normal, and isn't always an easy problem to solve. It's often easy to recognize the symptoms though, especially if you look at the trend over time. PingPlotter can really help pinpoint the problem, giving you the ability to see trending over time, latency patterns short and long-term, and latency / packet loss correlations.

6.10 Voice over IP (VoIP) troubleshooting

Introduction

Using an IP Network (like the internet) to conduct a voice conversation ([VoIP](#)) is becoming easier and easier for people to do. It can be inexpensive and relatively reliable. It can, however, also be challenging - with poor voice quality, the inability to hear and communicate, delays and other problems.

The underlying technology for VoIP is extremely network dependent. If you're having voice quality problems, the problem is often related to the network - maybe your internet provider or maybe some other component between you and the called party. This article will talk about some basic troubleshooting techniques that can be used to locate where the problem is occurring so you can make good decisions about how to solve the problem.

Network-related VoIP Symptoms

Many symptoms of VoIP problems are network related (although certainly not all of them). Here are some examples of symptoms that are often network related:

- Garbled words ("blips" and clicks mixed in with the words)
- Parts of words missing
- Gaps where the other side is talking, but you don't hear any of it
- High "distortion"
- Delays between the time you talk and the other side hears you (and vice versa)
- You start talking not realizing that the other party has started talking already too, and you talk over each other for a few seconds

Other symptoms might not be network related. In particular, if the symptom **always** happens, any time

of the day, any day of the week, then there's a decent chance it's not a network problem.

- Echo when you talk. This can be exacerbated by network problems, but constant echo is usually not caused by network problems.
- Inability to connect a call to some users while you can call others

Using PingPlotter to identify the source of network problems

PingPlotter has some unique capabilities in its ability to help you track down the source of network problems. What you really want to know is:

- Can **you** fix the problem or do you need to call someone else to fix it?
- If you need to call someone else, who do you call? Your ISP? Your VoIP provider?
- When you call them, how do you convince them it's their problem to fix?

PingPlotter can offer a **lot** of insight into all of these questions.

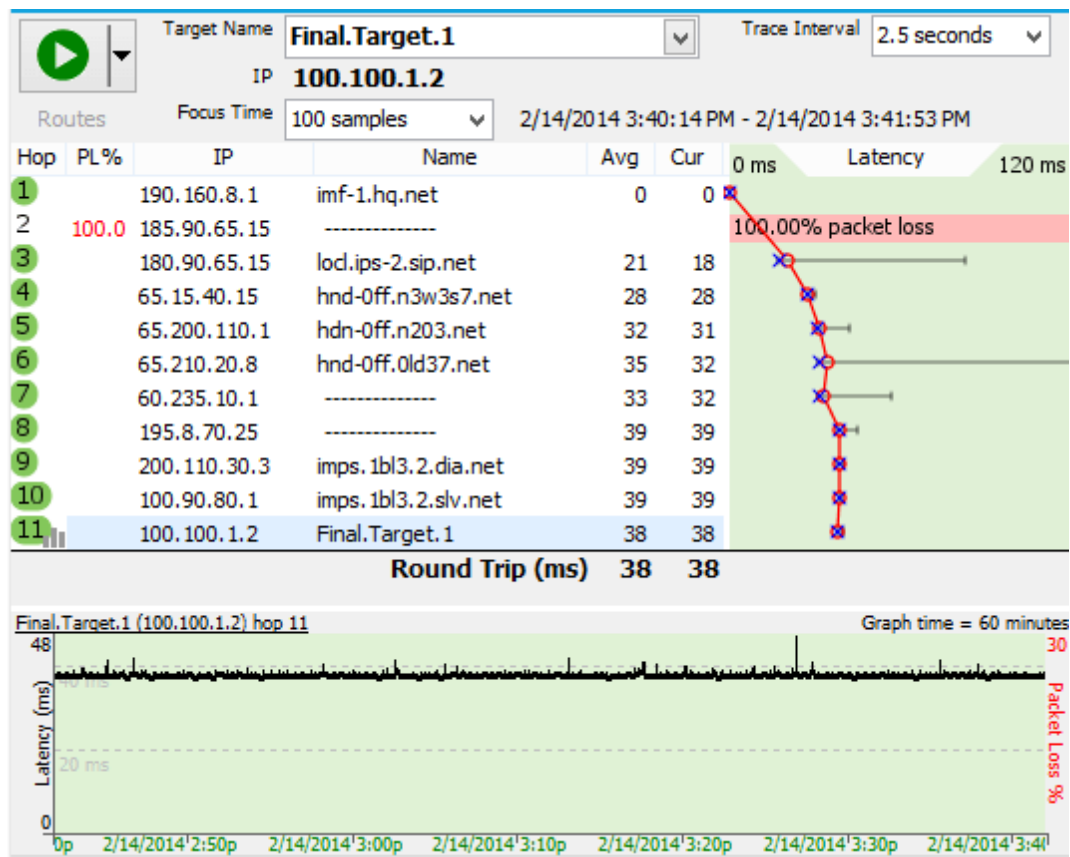
Collecting data with PingPlotter

Before we can do much analysis, we need some data to analyze. We cover some of these topics in earlier manual entries, so we won't cover *details* of how to set things up here.

First, we need a target server to monitor. Ideally, this would be the actual VoIP server of your VoIP provider, or something on the same area of the network. If you called your VoIP provider and they asked you to collect PingPlotter data, they may have given you a server to use. In many cases, the use of **any** server can work, but this will only identify problems with your ISP - not with your VoIP provider. The good news here, though, is that the vast majority of VoIP problems are because of front-line service providers (like your ISP). If you don't know what address to use and you have no way of finding out what address your VoIP hardware or software is using, try using the web site of your provider.

For this discussion, we will be a fictitious server (Final.Target.1) as the server we're monitoring and using for troubleshooting.

In your instance of PingPlotter, enter your VoIP server (or related target in the "Target name" field, and set the trace interval to 2.5 seconds. Now, hit the "Trace" button (or the "enter" key on your keyboard), and you should see a picture appear that looks something like this:



The upper graph should show a full route, including the "Round Trip". If you don't get a Round Trip, check in the troubleshooting section of this document for some ideas.

Now, let this run for at *least* 30 minutes - preferably, during a period where you're making a voice call. Ideally, you'll have a period where you have a voice call that's good and one that's bad, but that might be possible. If nothing else, just let it run for long enough to get a good sample of your network conditions.

A great thing to do while you're collecting data is to make notes in the PingPlotter data about what you're experiencing. You can see instructions on how to do this in our [Timeline Graphing](#) entry under the "Creating Comments" section.

The data we collected covers several days. PingPlotter works great to just run over a long period of time so you get a good idea of what network conditions look like - during good times and bad.

Examining data with PingPlotter

Once you've collected some data, it's time to have a look at what might be the problem. We cover some of the PingPlotter commands on zooming, focusing and digging in the [Finding the source of the problem](#) section.

One of the key things to know here is that we're looking for problems at the last hop only - and then using the other hops to determine where the problem starts. Packet loss or latency that [shows up only at an intermediate hop is not a problem!](#)

Let's look at the graph above. Notice how hop 2 has a full 100% packet loss? The final destination looks rock-solid, though - no packet loss and the latency, and it is mostly nice and smooth. This is, in general, what you want to see - a mostly flat line at the final destination, no packet loss (red lines in the time graph).

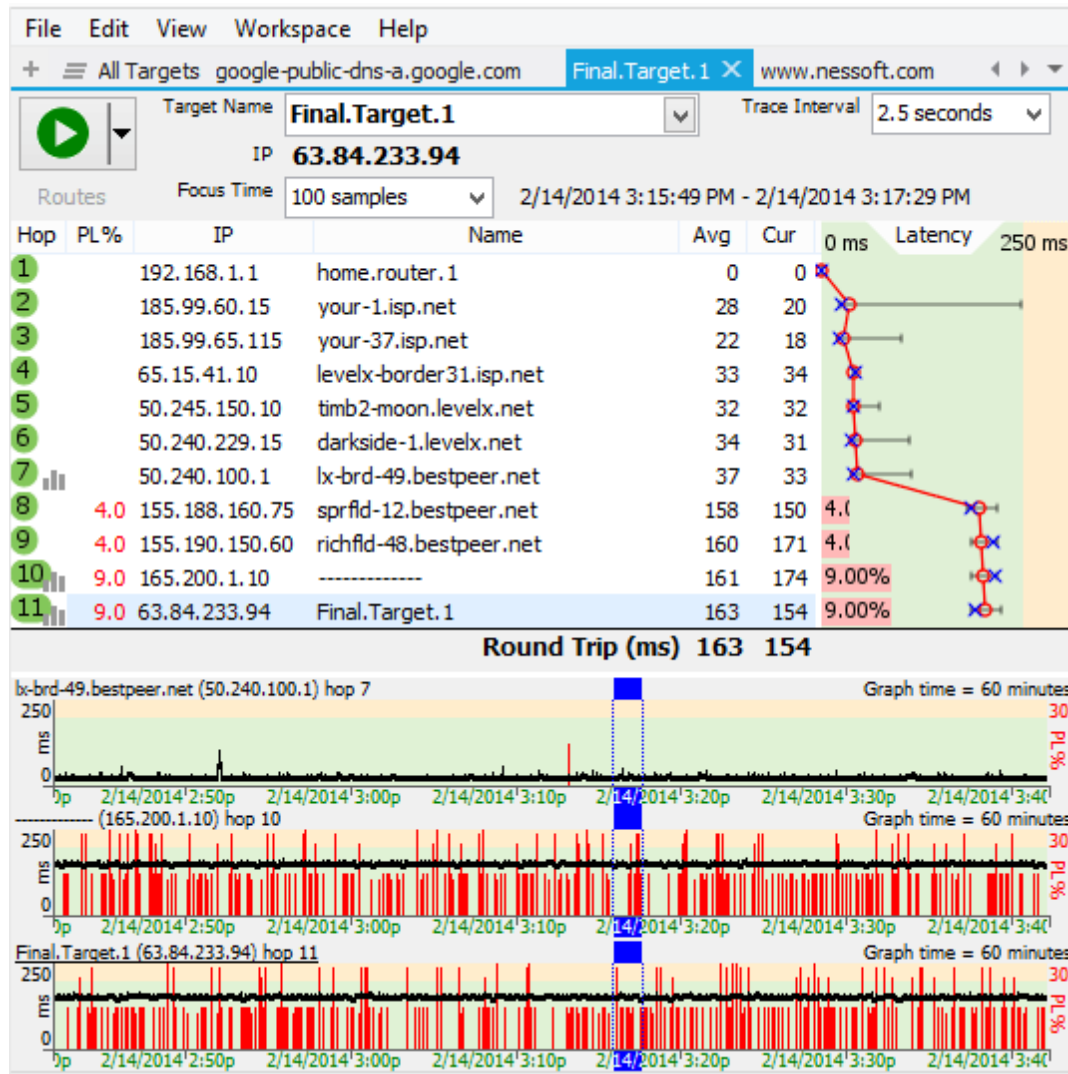
So, what are we looking for, when it comes to problems?

The first place to look is the final destination. If the time graph looks like the graph above (straight line, no red), then PingPlotter is not finding network problems. Look for problems at the final destination. If you find a problem at the final destination, then look back until you find the first hop showing similar symptoms - that's who we probably need to contact to get the problem corrected.

Examples and Analysis

Example: Distributed packet loss

Let's look at an example, this time an example with problems:



Here, we see 9% packet loss at hop 11 (the final destination). This would result in poor voice quality, dropped "bits" from words and hard to understand conversation. Notice that the latency is pretty good still - it's just the packet loss that's a problem (packet loss is all of the red in the time graphs and the red bars in the trace graph). With a pattern like this, voice quality would be consistently "iffy" - not unusable all the time, but not very good either.

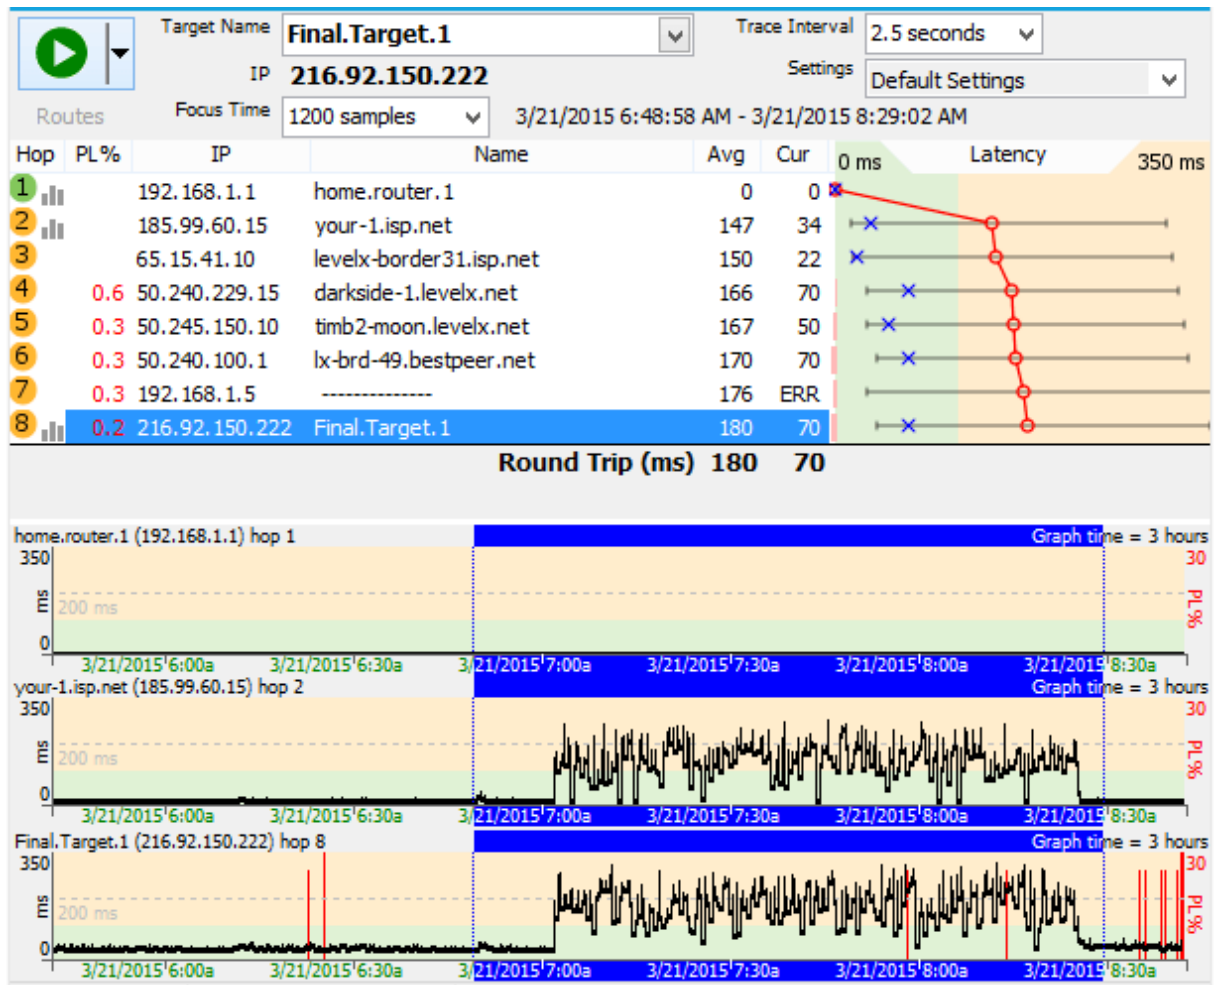
Notice how the packet loss is happening at all hops from hop 8 onward, while hop 7 looks relatively good. The packet loss percentage fairly similar all the way down (although, statistically, it would be just about impossible for all hops to have identical packet loss percentages with this kind of loss). To turn on and off time graphs like this, just double-click on the hop number in the trace graph.

So, in this situation, the problem looks to be between hop 7 and 8. It's pretty likely that Bestpeer knows about this problem - it's in the "middle" of their network - and it's all owned by Bestpeer (we can see that

from the DNS Name column).

In this case, we would need to contact Bestpeer about this problem. The picture above is pretty compelling and would be a good communication tool to them.

Example: Local bandwidth saturation



Here, notice the big latency jumps - you have a nice flat line, then a jump in latency, including some packet loss. This pattern is one that is almost always a bandwidth saturation issue (which is the same as congestion). In the case we have here, hop 1 is inside our network (our DSL modem, actually) and hop 2 is inside of our ISP.

This is a case where we were transferring too much during this period - and we were using all of our available bandwidth. A VoIP call would suffer significantly during these periods - there is a lot of jitter (the "ragged" line is an easy way to see jitter - where packets take different amounts of time to arrive), higher latency and some packet loss. The voice quality would be bad, there would be additional lag, and it

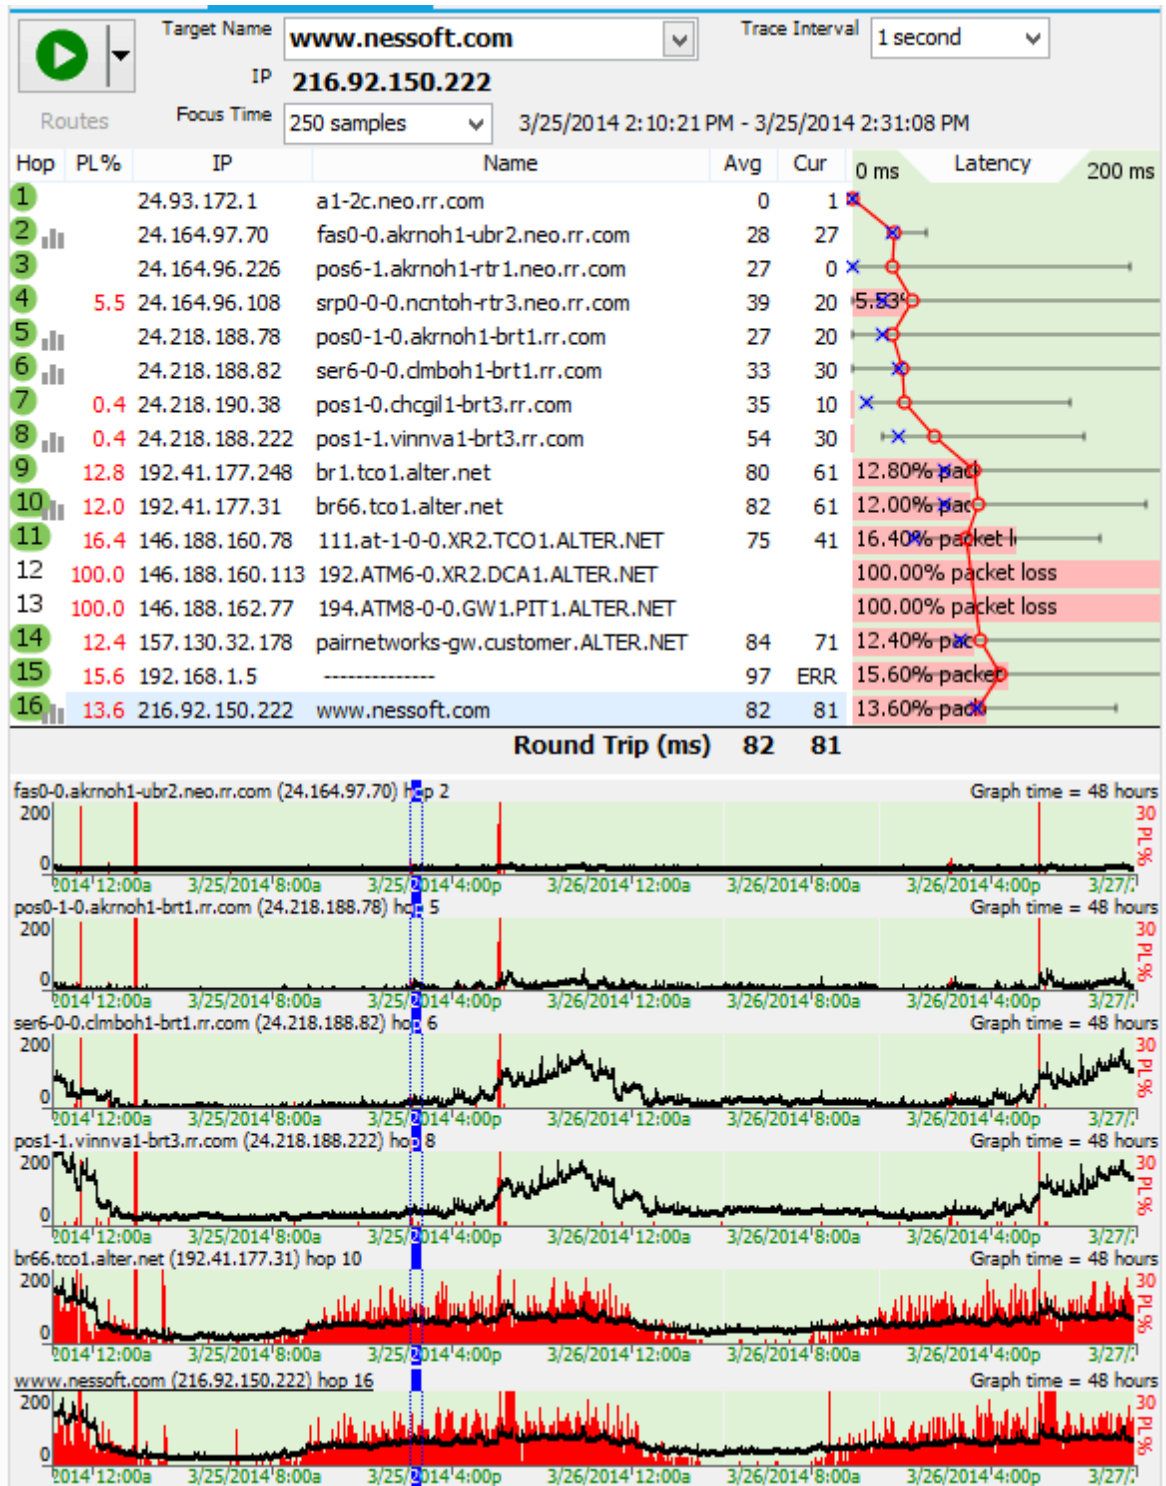
would probably have audio drops.

There are a few options for solving this one, but none of them involve complaining to anyone else:

- You can install a traffic shaping modem that gives higher priority to VoIP data (actually, unless you've configured PingPlotter packets to look like voice data, you might already have one of these in place - this article does not cover that topic, though).
- You can get more bandwidth (although that doesn't solve all problems - as you'll still be able to use all your bandwidth).
- You can use less bandwidth.
- You can get an additional broadband connection and dedicate it just to VoIP (this is an especially good idea for heavy VoIP users or businesses). The low cost of an additional broadband connection makes this viable in a lot of situations.

Example: Border congestion

Congestion often happens at network borders - where one network hands off to another. This is relatively common for small, growing ISPs - where they just do not subscribe to enough bandwidth to handle everything. Let's have a look at what this condition might look like.



This one isn't *quite* as simple, as there are a few factors. The symptoms of conditions like this would be:

- During high load times, this connection would be completely unusable for VoIP. With 14% packet loss

and really high jitter, this would be absolutely horrible for any kind of voice call.

- During the early mornings and late evening, it would be *bearable*, but the jitter would cause the words to be garbled sometimes, and not too fun.

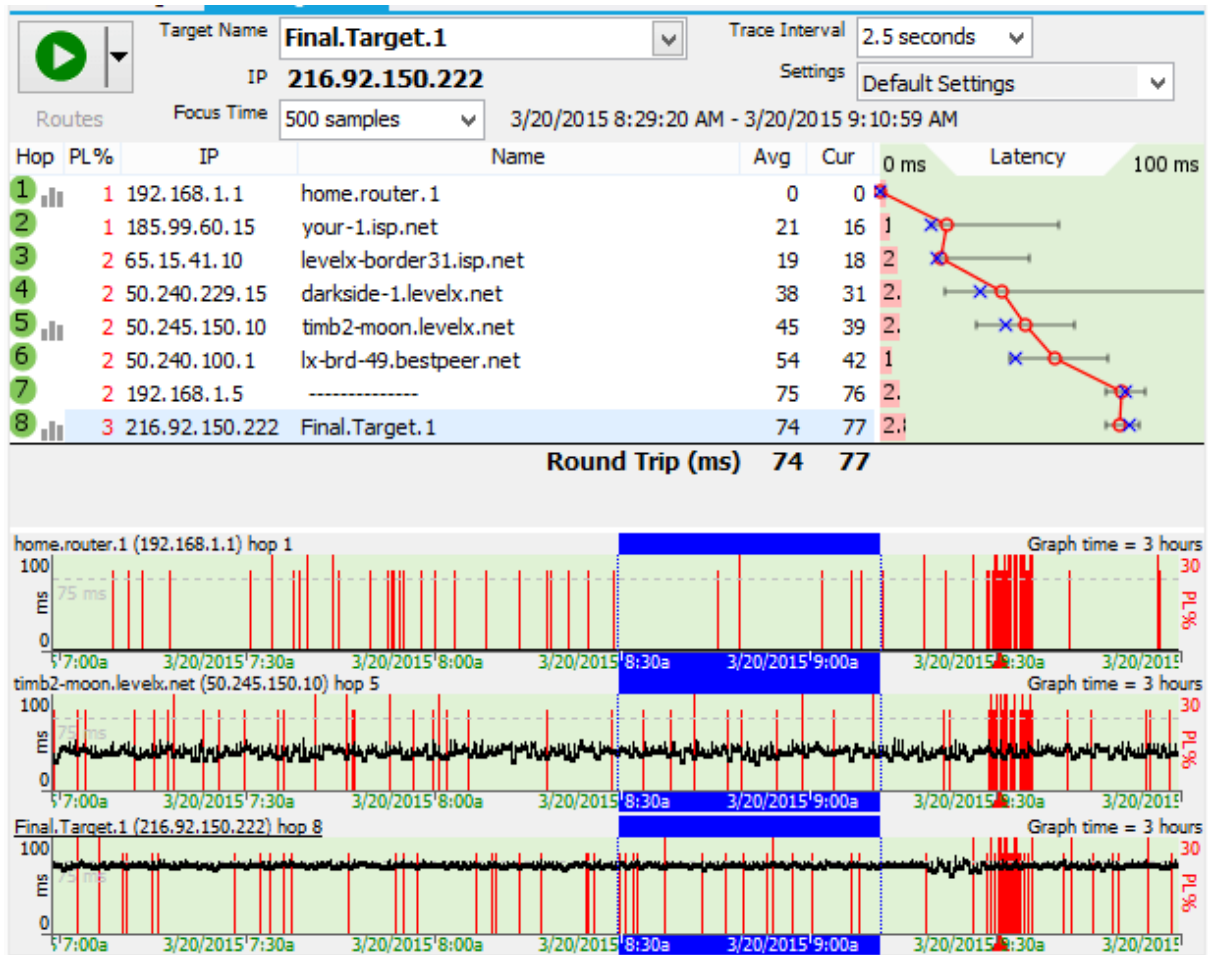
If we look at the network conditions with PingPlotter, we see a couple of problems. First off, there's some serious packet loss starting at hop 4. This packet loss is carried down through the rest of the route to the final destination. This is a border - between our ISP and levelx.net. Having problems at borders like this is pretty common - that's where one company pays another to handle traffic. If a company is growing, it might be "oversubscribed" - using more bandwidth than is available.

An interesting part of this is how during heavy load times for home users (ie: evening hours), the packet loss and latency are worse. During early morning hours, it goes back to being OK again. This is a big sign that the problem is load related - and that this link is having congestion problems at "rush hours". Time to add some lanes!

Another problem, though, is inside the rr.com network there is significant latency and jitter. There are some slight symptoms at hop 2 (which is the border between our internal network and rr.com - so that's the cable modem), but starting at hop 6 there's some real bumps in latency and the jitter (latency variation) is also a big cause for concern.

In this case, both problems are inside the rr.com network, and since we are an rr.com customer, we would call them for help on this.


Example: 802.11b network near its range limit



Here's an example where we're connected using a computer-based VoIP service (like Skype). Our computer is hooked up to our DSL modem via a wireless 802.11b network. Hop 1 is our DSL modem.

Here, we see a little bit of packet loss being added to every hop - our wireless network is losing a few packets (about 1 to 2 percent, it looks like), and this impacts everything this computer does - including our VoIP connection.

The call quality would be generally good here (probably better than acceptable - up to the "good" range, really). The latency is fine and there is very little jitter, but there is a little packet loss. There is a problem, though - at 9:31am, our call was interrupted - it looks like hop 1 lost a bunch of packets all together and during that period, we were unable to hear anything. Let's zoom in on that a bit.



See the period where we start getting a lot more packet loss, and then all the hops show a big block of lost packets - a period where it's likely no packets were getting through.

Here, the solution might be to move the wireless access point, or switch to wired on that computer.

Reporting problems, when you find them

If you're using PingPlotter, it's almost certainly because you're experiencing some kind of problem - and when you find something that you think might be the cause of that problem, you need to communicate that to the right party.

We [cover this topic](#) in some depth in our Getting Started Guide. The piece we want to stress here is that the data in PingPlotter doesn't really mean anything unless you correlate it with a network problem (like poor VoIP quality). It's of paramount importance that your complaints include a description of how this problem is affecting you. Don't just send a graph from PingPlotter expecting them to be able to figure out what was wrong.

One great way of doing this is to put comments in the PingPlotter graph itself using the "Create Comment" feature of the time graphs. Make comments every time your VoIP quality is bad. Make comments when you give up on a conversation because they can't hear you at all (but you can hear them just fine - how frustrating!).

Troubleshooting PingPlotter

If you get "Destination Unreachable" at something beyond hop 3 or so, but can access that site via a web browser.

Some sites do not respond to ICMP echo requests. See our knowledge base article for [instructions on how to configure PingPlotter to use TCP packets](#) instead of ICMP.

If you get "Destination Unreachable" at hop 1 for all targets

Make sure your software firewall (ie: ZoneAlarm, etc) is configured to allow PingPlotter to have access to the network.


If you only have the final hop visible - and all intermediate hops are empty

We cover this in this [knowledgebase article](#).

Other questions

Jitter

Jitter is the amount of variation in latency. If one packet takes 100ms and the next one takes 200ms, there's 100ms of jitter there. PingPlotter Pro offers jitter calculations and graphs, but PingPlotter Standard (and the 2.x line) still gives you an easy way to see the jitter by looking at the smoothness of the time graph, zoomed in a little. Here's an example:



This is zoomed in enough for us to see the individual samples, and we can see that none of them come in with the same latency. Adjacent samples here often have latency variations of 100ms, and just about every one has latency variation of 30ms or higher. Just looking at this graph, we can see a lot of jitter. Compare that to the first picture in this article - where the line was completely flat. We're looking for the flat lines, not big variations with red stuck in everywhere.

Other resources

This article introduces some concepts and ideas about VoIP troubleshooting. There are other resources online that provide more depth (albeit not within the context of PingPlotter).

www.whichvoip.com/voip-troubleshooting.htm - a great page with real-world solutions and suggestions.

Targeted to residential, but useful everywhere.

www.voiptroubleshooter.com is a great site that has an enormous amount of content on which symptoms relate to what kinds of network problems. This site has a relationship with [Telchemy](#), a leading VoIP provider of call quality monitoring tools.

6.11 Building a Compelling Case

If you find a network problem, you'll most likely want to try and solve that problem. In some cases you might be able to solve this network problem yourself - possibly by upgrading the BIOS on a hardware device, replacing a network cable or changing network service providers. In most cases though, the network problem will need to be fixed by someone else. If this happens, you'll want to build a compelling case that clearly demonstrates the problem, and then also convey (or present) that information. PingPlotter can help you do both of these tasks. It does this by collecting data over time, and then giving you the capability to present it in a way that can be compelling to someone else.

In this section, we'll show you how to use some of the tools in PingPlotter to collect data, save and reload that data, and then cover areas to focus on to make a really convincing story.

A convincing story (when using PingPlotter to document a network problem) is one that has some (or most) of the following attributes:

- Correlates the problem description (i.e.: bad VoIP quality or slow game performance) with the supporting data (i.e.: packet loss in PingPlotter).
- Is clear about the actual problem being experienced. Bad PingPlotter data is meaningless unless there is some impact to other applications. Make sure you describe the affect this problem is having on your network experience.
- Includes solid supporting data, covering at least a few minutes and possibly several days
- Does not exaggerate the problem or the data. You can certainly zoom PingPlotter in to only the very worst 10 samples you collected, but that doesn't give a realistic picture of the problem
- Is concise and not argumentative.

Most of the time your first contact with your provider (the one who can help you solve the problem) is going to be with a first-level, front-line support technician. In many cases, their goal is to "close the case". You need to be polite but persistent, and have a strong story to get them to bring the problem to the next level.

Because many ISPs and providers don't give copies of PingPlotter to their front-line support technicians, you probably won't be able to just send a PingPlotter save file for them to analyze. You'll need to create an image they can easily see in their email, or attach to their support case.

Creating a graphical image of the problem

The main graphs in PingPlotter can be a compelling picture of the problem. In a lot of cases, you'll want

to include a snapshot of this screen in an email. There are a few ways to do this.

- You can use the ["Send Email" option under the File menu](#)^[122] to send an image via email.
- Depending on your email software, you might find it easiest to just paste an image into an email or document. To do this, use the ["Copy as Image" option under the Edit menu](#)^[123]. This will put the image into the clipboard, and you can then paste it into your email or other document.
- You might, instead, prefer to create an image file and then attach it to an email. To do this, use the ["Save Image..." option under the File menu](#)^[122]. This will allow you to save as a Windows bitmap, GIF or PNG file. Out of the box, the default file type is PNG, which creates the smallest possible file (with the highest quality as well!).

PingPlotter also supports automatic saving of image files. See [our documentation on this topic](#)^[119] for more details.

Part



Advanced Features

7 Advanced Features

7.1 Raw TCP sockets and WinPcap

Starting with PingPlotter 2.60, TCP packets can be used in addition to the previously available ICMP and UDP packet types. TCP traceroute allows access to some targets that were not previously available, but this comes with some caveats.

One challenge that many users will face is that Windows started blocking the ability to create TCP packets with the options needed for traceroute (starting with Windows XP SP2). See [a Microsoft white paper](#) for some more details on this.

There are several ways to "work around" this limitation in Windows, including:

- Turn off the Windows firewall service. Issuing a "net stop sharedaccess" command from a command prompt will disable the raw socket block.
- A piece of software (ie: PingPlotter) can use a device driver to bypass this restriction.

Since many of our users rely on Windows firewall service and/or Windows ICS (internet connection sharing), we've implemented a solution that uses a device driver to create TCP raw sockets.

Do you need to use WinPcap?

If you want to use TCP traceroute and your operating system is Windows XP SP2 or newer, then you will need to use WinPCap.

If you're using an older operating system (Windows 98, or Windows 2000 for example), then you probably don't need to use WinPcap. We'll update this page with more information if we find other situations where it's helpful.

How to install the driver

To make this as reliable as we could, we decided to use an open source (and free) driver, [WinPcap](#) to send raw sockets. This driver needs to be installed before PingPlotter can use it, and should be downloaded and installed based on the WinPcap instructions.

Install steps:

- Visit the [WinPcap site](#) so you know what you're installing.
- [Download the driver](#). We recommend using the latest version. If you're using Windows Vista x64, then 4.0 or higher is required.

- Install the driver. The [WinPcap download page](#) has reasonably good instructions. Basically, run the installer.
- Make sure PingPlotter is configured to use WinPcap (the default configuration is to use it if its installed, so it should "just work").
- You're done!

Shortcomings with WinPcap and PingPlotter

WinPcap allows PingPlotter to send packets directly to the network card. This is very powerful, but also means that the Windows protocol stacks don't help us with routing or validation of the packet.

PingPlotter should do a good job with detecting if a gateway should be used and sending to that gateway, but might not work correctly with multiple gateways, or other complicated network setups.

Please [contact us](#) if you find a network environment where PingPlotter seems to be making poor decisions or causing problems.

If you're not running as administrator (if you're using Windows Vista, for example), then you'll need to either manually start the WinPCap driver or run PingPlotter as administrator. For more details on how to do this, see the WinPCap FAQ, [question 7](#) and [question 18](#).

7.2 Command Line Arguments

You can have PingPlotter do a few things automatically on startup by specifying command line parameters. You can put these parameters in a shortcut - or enter them from a DOS command line window.

```
PingPlotter.exe [File to Load] [/TRACE:[Address To Trace]] [/SAVE] [/INIFILE:(filename)]
```

Loading a file at startup

If you enter a parameter without a / on it, PingPlotter will try and find a file by this name - and load it if it's found.

Example:

```
PingPlotter.exe www.pingplotter.com.pp2
```

/TRACE

This option will start tracing automatically when PingPlotter loads. If use the option to load a save file on

startup, then tracing will begin to this address. Otherwise, Add a colon (:) and the IP Address or server name you want to trace to.

Example:

```
PingPlotter.exe www.pingplotter.com.pp2 /TRACE
```

or

```
PingPlotter.exe /TRACE:www.pingplotter.com
```

/SAVE

This option can only be used when you specify a file name on the command line. If you use /SAVE, then any new traces (to the original address specified in the save file only) will be saved to that file on shutdown automatically (without asking you if you want to save).

Example:

```
PingPlotter.exe www.pingplotter.com.pp2 /TRACE /SAVE
```

/INIFILE

You may want to start PingPlotter with a different set of parameters and setups. This is particularly useful if you're auto-starting PingPlotter and having it trace automatically. In this situation, you may want to have different trace intervals, or graph times, or whatever. Starting PingPlotter with an alternate INI file allows you to save multiple setups and use these different setups as needed.

Example:

```
PingPlotter.exe /INIFILE:alternatesetup.ini
```

/?

Show this help screen.

Example:

```
PingPlotter.exe ?
```

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

7.3 Custom time graph intervals

You can change, and also add, additional timeline intervals for the timeline graph if the defaults available from the right-click menu don't fit your needs. For example, you may want to look at results farther out

than the default maximum value.

The default set isn't written in the .INI file, however you can add an entry and it will override the default values.

For example, if we wanted to add "2 hours" and "14 days" intervals to the right-click menu, your .INI file entry would be what you see in the example below. It's important to note that the values are in minutes.

It's also important to note the "Count" value. If the value was set to "Count=12" in the example below, the value for "Interval13" wouldn't show up in the menu.

```
[TimeGraphIntervals]
Count=13
Interval1=1
Interval2=5
Interval3=10
Interval4=30
Interval5=60
Interval6=120
Interval7=180
Interval8=360
Interval9=720
Interval10=1440
Interval11=2880
Interval12=10080
Interval13=20160
```

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

7.4 Installer Options and MSI

The PingPlotter install is MSI-based, and wrapped with a bootstrap that helps do upgrades.

Normally, the best way to install PingPlotter is just to launch the installer by downloading and double-clicking.

In some cases, though (particularly for deployment to multiple computers), it may be helpful to change the way PingPlotter is installed.

The bootstrap has several options to extract the MSI, log the install to file, debug, etc. To see the

options, launch the installer from a command line and pass it a `/?` parameter - this will list the parameters that can be used.

License Entry

PingPlotter license are stored in the registry in the following location:

```
[HKEY_LOCAL_MACHINE\Software\Pingman Tools\PingPlotter\User]
"UserName"="Your Username"
"RegistrationCode"="Your License Key"
```

If you need to automate license key entry, you can write these values into the registry.

*** 64 bit windows warning *** - Since PingPlotter is 32 bit application, on a 64 bit machine, this needs to be written to:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Pingman Tools\PingPlotter\User]
```

Custom Options

There are several other options available through the command line interface of the installer. Please [contact](#) our support team with your needs, and we'll help you build the right install package.

7.5 Whois Setup

If you'd like to add additional whois servers to have available from the trace graph's right-click menu, you can edit the pingplotter.ini file. Currently PingPlotter only supports xxx.yyy style addresses. Not, for example, xxx.yyy.uk. You can also change the default servers directly in the .ini file instead of changing them from the Options/Internet menu.

To change the default settings, add the following to your pingplotter.ini file:

```
[Internet]
StandardWhoisHost=whois.networksolutions.com
BlockIPWhoisHost=whois.arin.net
```

Also, you can add additional servers with a list of commands like the following:

```
[Internet]
AddlWhoIsServers=Internic Whois Lookup, whois.networksolutions.com, name,
Arin (IP) Whois Lookup, whois.arin.net, ip
```

Basically, each whois server is defined by 3 settings: **Description** (as it appears on the menu), **Server address** (IP or name), and **lookup type** (Name or IP). The last setting specifies if it queries the whois server for the IP address, or the name of the specific hop.

If you're adding a server that's not specified in PingPlotter by default, add that server as the first server on the AddIWholsServers line (i.e.: don't include the default servers shown in the example above, since they're always included no matter what you put in the INI file). If you have more than 1, just add 3 more comma delimited sections. There is no limit to the number of sections you can include, though more than a few will make the menu a bit unwieldy.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

7.6 Running from a USB drive

PingPlotter can be run from a removable drive (like a USB key) without having to be installed, or modifying anything on the "host" machine where you're running PingPlotter. This is great if you administer a number of PCs and want to be able to run PingPlotter without having to install it, and/or enter a license key.

There are several differences in operation when PingPlotter is running from a USB drive in special USB mode.

- The license key is stored on the USB drive, not in the system registry. This means you don't have to enter the license key on each machine you're running on.
- All settings that would *normally* change the registry will no longer work. For example, you cannot install as a service when running from a USB drive. Also, the .pp2 file association is not tied to PingPlotter automatically.
- All PingPlotter settings are written to the .ini file on the USB key (in .exe directory - so make sure it's writable!).
- PingPlotter does not check in the user's profile for override settings.

To run PingPlotter in this special mode, follow these steps:


1. Install PingPlotter normally on your computer's hard drive (not the USB drive).
2. Copy the entire directory where PingPlotter is installed to your USB drive. This should include all subdirectories that may exist. You can put this in any directory on the USB drive, but we'll assume

this is in f:\PingPlotter for now. The executable to launch PingPlotter would then be f:\PingPlotter\PingPlotter.exe

3. Uninstall PingPlotter from your computer's hard drive.
4. Create an empty file named "license.dat" into the directory where PingPlotter.exe exists on the USB drive (in our example, this is f:\PingPlotter).
5. Launch PingPlotter from the USB drive. When prompted, enter your license key. PingPlotter will put your license information into the license.dat file.
6. Create an Autorun.inf (outside this scope of this topic, but email us if you need help) or hook up the f:\PingPlotter\PingPlotter.exe file to your USB drive's menu system.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

Part



Reference

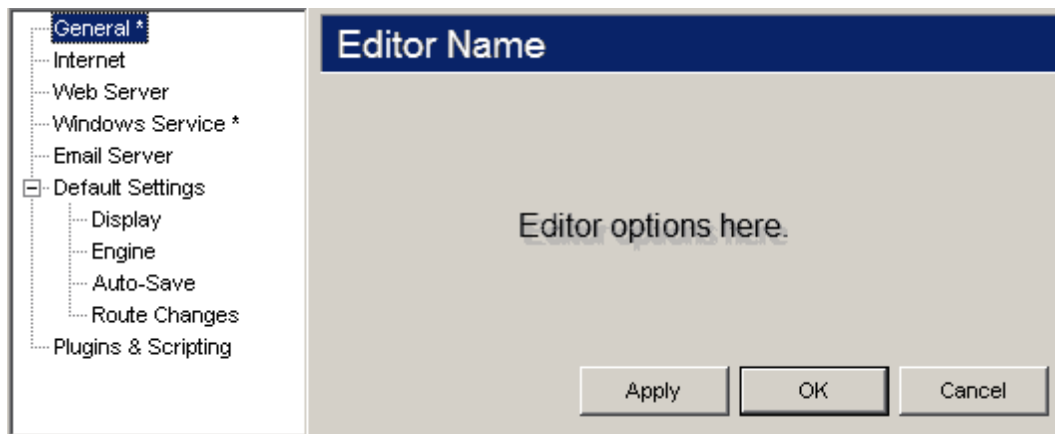
8 Reference

8.1 Options

The options dialog is the main way to configure PingPlotter. There are a few interesting behaviors here that may not be obvious.

Applying (and identifying) Pending Changes

First, any changes made to these settings do not take affect until you either hit the "Apply" or "OK" button. The Apply button will become enabled when there are any changes pending to be applied.



The topics in the tree view with the * marks (asterisk) are the ones that have changed. Hitting the Apply key will write these changes to PingPlotter and put them into affect. Hitting the "OK" button will do the same, but will also close the options dialog. Hitting cancel will back out the changes that have not been applied. You can make changes to multiple areas before applying them.

Creating / Manipulating named configurations

PingPlotter Pro supports multiple named configurations. Right clicking the "Default Settings" option above will allow you to create more.

If you delete a named configuration that is currently in use by a target, that target will automatically change to use the first configuration in the list.

Making "on the fly" changes while still using PingPlotter

If you want to make changes in the options dialog will still interacting with the main PingPlottergram, hold the control key down while you launch the options dialog. The options dialog will stay on top, but it will allow you to interact with other areas of PingPlotter at the same time. As soon as you hit the "Apply"

button, these changes will go into affect - without you having to close the options dialog.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.1.1 General Options

General Settings

General Options

- Put icon in tool tray?
- Show "Round Trip Time" row?
- Prompt to save on exit if lots of new data exists
- Minimize PingPlotter when Windows "close" command used (ie: X on upper right)
- Include settings name in target / host descriptions (tabs, time graphs, etc.)

Show welcome splash on startup: 10 seconds

Summary Graph

- Automatically show final destination on default summary screen.
- Automatically show timeline graphs for targets added to summary screen.

Workspace

- Save collected data with workspaces.
- Auto-save/load active workspace when PingPlotter is closed/opened.
- Auto-save active workspace every: 30 minutes

Put icon in tool tray?

PingPlotter can be minimized to the "tool tray", the small icon "tray" where the clock normally sits, and where a number of other notification icons might appear.

Enabling this option will turn on a tray icon at all times. When PingPlotter is minimized, it will only be visible in the tray. When it's not minimized, it will show on the taskbar and the tool tray. Alert conditions can be surfaced through the tool tray as well, in which case the icon will change to red and a message might appear.

Show "Round Trip Time" row?


The "Round Trip" row of PingPlotter duplicates the information from the final hop and makes it evident that the host is reachable and what the round trip time and latency is. Before adding this option, we got a lot of questions "What's the round trip time?"

Prompt to save on exit if lots of new data exists

When PingPlotter is closed and has been collecting data, you might not want to just throw away all that data. If this checkmark is on and there are at least 75 unsaved samples in memory, PingPlotter will ask if you want to save it when closing.

Minimize PingPlotter when Windows "close" command is used

Turn on this option to cause PingPlotter to minimize instead of close when the "X" button is hit.

If you normally run PingPlotter all the time, you might not want it to close if you accidentally hit the "close" button on the application (ie: the  button). Turning on this option will make PingPlotter minimize instead of close. To close PingPotter, use the "File" -> "Exit" menu option, or if PingPlotter is minimized to the tray, use the right-click "Close" command. Note: using the close command from the taskbar will *not* close PingPlotter, as that is equivalent to using the X button.

Include settings name in target / host descriptions

In PingPlotter Pro, when tracing to the same target with different engine settings (see our named configurations documentation for more details), the only way to distinguish between the settings is by the named configuration. If you're only using one named configuration (or if you're tracing to all different targets), then this is not so helpful. Turn on this option to show the named configuration on all tabs and time graphs. For the summary graph, turn on the "Settings" column to show the named configuration used for that target.

Show Welcome Splash on Startup

This option allows you to adjust the amount of time the splash screen is shown on start up of the program (there's also an option in the drop-down to not show the splash screen at all).

Summary Graph Settings

Summary graphs are exclusive to PingPlotter Pro.

Automatically show final destination on default summary screen

This is on by default, but can be turned off if you want to manually add your own targets to the summary screen, rather than have PingPlotter automatically add them for you.

If this option is on, when you start a new target in PingPlotter, that target will automatically add itself to the summary graph screen. The summary graph screen is a handy way to see the current status.

Automatically show timeline graphs for targets added to summary screen

If you regularly trace to a lot of targets (and have them auto-show on the summary screen), turn this option off to control visibility manually.

When a new target or router is added to the summary screen, having a time graph automatically show up can be handy. The downside of this is that a long list of targets quickly fills up the screen and becomes less useful. Turn off this option if you find yourself regularly turning off a lot of the automatically added graphs. You can turn them back on at any time manually anyway.

Workspace Settings

Workspaces are exclusive to PingPlotter Pro.

Save collected data with workspaces.

A handy option that is usually turned on - this makes it easy to save a workspace and reload with everything just the way it was. If disabled, reloading the workspace will *not* reload the collected data, just the targets and layout. The data will be saved in a directory named the same as your workspace (without the file extension).

Auto-save/load active workspace when PingPlotter is closed/opened.

A great way to have PingPlotter remember everything you were working on and automatically pick it back up when you reload. When you close PingPlotter, everything is saved. When you re-open, it's loaded again.

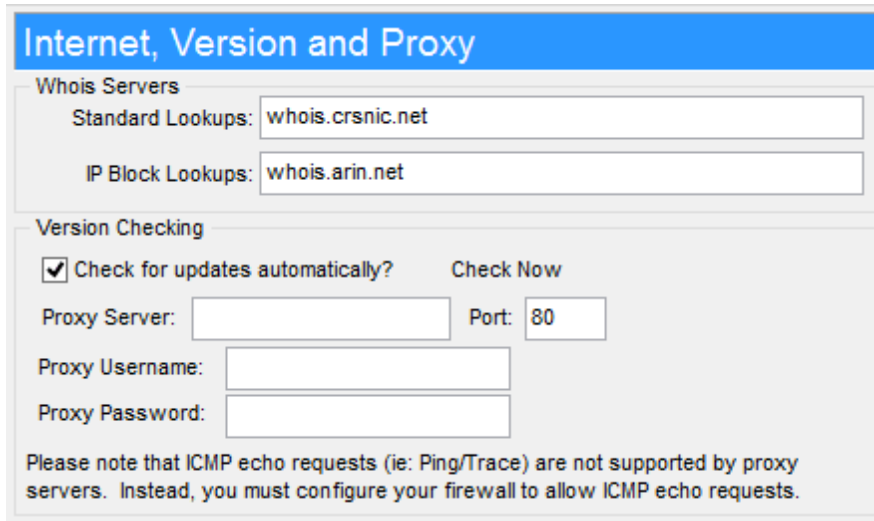
If you have this on and want to clear everything and start fresh, you can do that via the workspace menu.

Auto-save active workspace every:

Manually key a time into this field, like "30 minutes" or "6 hours". Click out or tab out of the field to "verify" what you entered and have PingPlotter interpret it. 30-60 minutes is a pretty good number to use here. Numbers like "1 second" mean that your computer will be constantly saving things, and probably won't have much free time for anything else - especially if you get a lot of data collected.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.1.2 Internet Options



Internet, Version and Proxy

Whois Servers

Standard Lookups:

IP Block Lookups:

Version Checking

Check for updates automatically?

Proxy Server: Port:

Proxy Username:

Proxy Password:

Please note that ICMP echo requests (ie: Ping/Trace) are not supported by proxy servers. Instead, you must configure your firewall to allow ICMP echo requests.

The "Internet" settings control the connectivity for whois, update checks, etc.

Standard Whois Lookups

Right-clicking a hop on the upper "route" graph on any hop that has a DNS name defined gives us the option of looking up that name and seeing who the owner is.

The "Standard Lookups" address is the WHOIS server to query for named lookups. The default value is whois.crsnic.net. This whois server is a central area which doesn't actually give whois information directly, instead it tells us which domain registrar handles the selected domain. PingPlotter then uses this information to go to the registrar and do a whois lookup again. Note that whois.networksolutions.com is another whois server that works - and will tell which registrar is responsible for a domain in most cases.

IP Block Lookups

Internet IP Addresses are assigned ownership, and sometimes its interesting to find out who owns them. The "IP Block Lookups" address is the WHOIS server to query for IP block lookups. This will look up who owns a particular IP address. The default for this is whois.arin.net.

Version Checking

PingPlotter can check with the PingPlotter servers occasionally to see if there's a new version available. This can be done automatically, or it can be done manually by hitting the "Check Now" button.

Proxy Setup

If you access the internet through a proxy server, you can set up this server here. In the current version

of PingPlotter, this is not required, although automatic version checking doesn't work if PingPlotter can't access the internet via HTTP.

The Proxy Setup is currently used for only one thing – doing the version check. PingPlotter doesn't use these settings to do the trace in any way.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.1.3 Email Settings

The email setup dialog is used to set up emailing for alerts. If you're not using alerts, or you're not interested in having the alert system email you, then setting this up is not required.

Email Server and Identity

These settings are used for sending alert emails. All outgoing alert emails will use these configuration values.

Address and Options

Return Address:

SMTP Server

SMTP Server:

Server Port (blank = 25):

Use SMTP Authentication?

SMTP Server Username:

SMTP Server Password:

Use STARTTLS encryption if server supports it.

Return Address

All outgoing emails will have a return address specified, and this is the address that is used. Please make sure you specify a valid address here since this is where all the bounce messages will come from. Some ISP SMTP servers only allow emails sent out with a "from" address of their domain as well, so if you're having problems getting the SMTP server to work, make sure you're using a valid return address.

SMTP Server

The SMTP server is the server that your outgoing mails will go through. This may have been given to you by your ISP or your mail administrator.

Server Port

The default port for most SMTP servers is 25. If you connect to your SMTP server via a different port, then enter that port here. Leaving this blank will use port 25. If you're using STARTTLS/SSL, then this might be port 587 or some other port as supplied by your email server/provider.

SMTP Authentication

Some SMTP servers require a username and password to be able to deliver mail. If this is the case with your server, turn on the "Use SMTP Authentication" checkbox, and then enter your username and password. The password is saved in your PingPlotter.ini file using a basic XOR encryption scheme – this will keep your password hidden, but this encryption method is "crackable" if someone really wants to figure it out by looking at your .ini file and reverse engineering it.

Using STARTTLS encryption

PingPlotter also supports use of SSL (STARTTLS) for SMTP. For more details on this, [see our web site](#).

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.1.4 Display Options

Default Settings: Display Settings

Color HOP column by speed

Color graph background

Warning (yellow) speed (ms):

Critical (red) speed (ms):

Draw X on graph for current sample

Draw line to show Min/Max range

Graph Scale

Automatically Scale to Sample Times

Scale to: ms

Packet Loss

Don't graph packet loss

Packet Loss Scale at %

Jitter

Add jitter graph above time graphs?

Jitter Graph Scale: ms Jitter "Target" line: ms

Graph Height: pixels Hide if area less than pixels

The "Display" settings control the general display format of PingPlotter's graphs, including scaling, coloring, and other general values.

Color HOP column by speed

If you want to minimize on-screen color, turning this off will hide the green, yellow and red background color on the hop (first) column.

Color graph background

The background of the graph uses colors that don't display well unless your video drivers are set for more than 256 color. Turn off this option if you're having problems seeing the graph (ie: It has little speckled dots instead of a solid color).

Some 256 color graphics drivers might not look very good with the colors on the graph. Turn off this switch to use a non-color background.

Warning and Critical speed limits

These boxes control the point at which the colors change. By default, all response 200 ms and below will paint green. From 201 to 500 will paint yellow, and over 500 will paint red. These numbers apply to both the HOP column and the graph background. In addition, the legend on the graph screen will be updated with these number.

You'll probably want to change the numbers based on your internet connection speed. If you've got a T1 or a cable modem, the listed numbers are probably pretty good (you might move them down a little if you're tracing to a fast site). If you have a modem, you probably want to crank these numbers up a bit.

Changing these values sets the Green / Yellow / Red threshold for the graphs. This is dependent on your expected performance. For a modem, 200 ms might be quite good, while for a T1, it could be considered bad.

Draw X on graph for current sample

To watch "trending", it's sometimes nice to see what the most recent sample is. This option will enable that. A little blue X will be drawn on the graph that represents the most recent sample. You might want to turn this off when submitting a picture to an ISP so as not to confuse them or add anything they can question.

Draw line to show Min/Max Range

Show the Min/Max Line. Hide this line to keep the scale of the upper graph in better range. This line can be useful to understand how a specific hop is responding - for example, if hop 8's minimum point is significantly greater than hop 7's maximum point, then you may need to investigate what's happening

between hops 7 and 8. It may be distance (ie: speed of light latency), or it may be a problem with one router, or the connection between those routers.

When showing just a few samples, this can be really handy to see the range of latencies. As you increase your window, though, a single bad sample can make this line stretch the scale of the graph.

Graph Scale

The number shown here indicates the graph scale, in milliseconds (1/1000th of a second).

If a dynamic scale is being used, then this number is the maximum response time of any of the included sample set. This number can change (and WILL change) as new samples are received.

If a fixed scale is being used, this number will always equal that scale. You can change to a fixed scale in the options screen.

Normally, setting the graph scale to automatic works pretty well. Sometimes, you might get a few samples *way* out of range, though, that stretches the scale. This is especially likely as you increase the value for samples to include.

Packet Loss

The red number on the right of the timeline graph is the scale of the packet loss numbers. Depending on the number of samples included in the timeline graph, all timeouts may show 100%. For more details, [click here](#).

Most often, graphing the packet loss is a handy, easy way to see lost samples. 30% seems to work great for highlighting just the right of loss in most cases, but you're certainly going to run into cases where you want to change this to something lower (as low as 1) or higher (any number is valid – even over 100).

Jitter

Jitter is a number that represents how stable the latency responses have been. A low jitter number is usually an indicator of a good connection. High jitter can lead to slow response times, poor voice quality (in Voice over IP) and other connection problems.

PingPlotter Pro allows you to graph jitter correlated with the time graph. In many cases, jitter is apparent when examining the standard PingPlotter time graphs, so the jitter graph is only displayed when there is enough room. The settings here allows you to control when that is displayed, and what it will look like.

Jitter Graph Scale and Target line

The jitter graph scale controls the range of jitters you expect, and at what point the jitter will go off-scale.

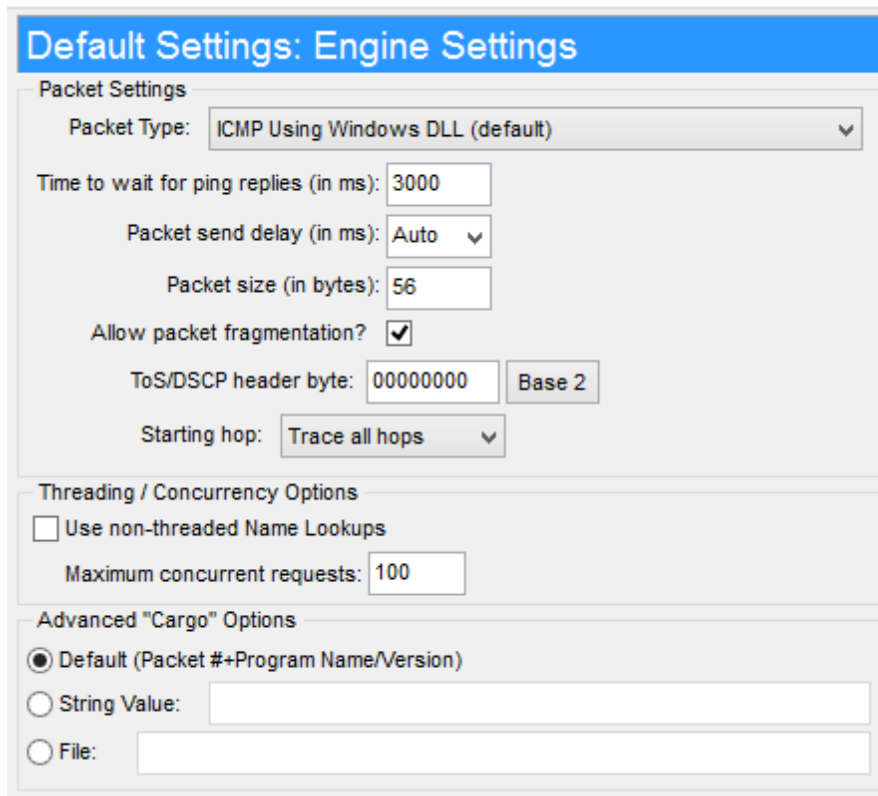
Normally, a jitter of much over 60 ms is indicative of a problem, so 60 is a good starting point. Any values over the scale will show with a red line for "overscale".

This graph scale is used in tandem with a "target line" which will be drawn across the graph at the point you specify. This is used to easily identify if jitter is exceeding your target number.

These settings are represented in milliseconds, and control the range of jitter values that will fit into the graph. Note that this also applies to the web interface, so if you want the jitter graphs to show up in the web interface, these settings need to fit the height of the web interface time graphs.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.1.5 Engine Options



Default Settings: Engine Settings

Packet Settings

Packet Type: ICMP Using Windows DLL (default)

Time to wait for ping replies (in ms): 3000

Packet send delay (in ms): Auto

Packet size (in bytes): 56

Allow packet fragmentation?

ToS/DSCP header byte: 00000000 Base 2

Starting hop: Trace all hops

Threading / Concurrency Options

Use non-threaded Name Lookups

Maximum concurrent requests: 100

Advanced "Cargo" Options

Default (Packet #+Program Name/Version)

String Value:

File:

The "Engine" settings control what and how PingPlotter sends data.

Packet Type

The "Packet Type" settings allows you to pick what kind of data you want to send to tailor PingPlotter to your network needs. PingPlotter supports 4 packet types:

- **ICMP using Windows DLL.** This method is the traditional method and matches the data that the Windows TRACERT command uses. It works on all Windows operating systems, and is a good balance of reliability and capability. This is reliable with the least CPU usage (on most operating systems). This method will automatically do manual timings on less-than-accurate operating systems to attempt to get accuracy to 1ms. This method does not require administrative rights, and should be the first choice for most users.
- **ICMP using Raw Sockets (advanced use only).** In some rare cases, the standard Windows method doesn't work. PingPlotter can compose its own ICMP packets - although in most cases this is no more reliable or better than ICMP.DLL. This requires administrative rights, and doesn't work reliably on Windows Vista or newer (including Windows 8).
- **UDP Packets (Unix-Style).** This uses ports 33434 – 33500 and closely mirrors Unix's traceroute command. This method will sometimes allow you to trace to a destination that isn't reachable via ICMP, or might allow you to reach the internet even if your ISP is blocking ICMP echo requests. Though not the cure-all for "Destination Unreachable" issues, this is worth a try, especially if you're getting erratic packet loss or unreachable destination. This requires administrative rights, and doesn't work reliably on Windows Vista or newer (including Windows 8).
- **TCP Packets.** This method gives you the opportunity to send TCP packets. If a firewall is blocking ICMP packets, it's sometimes possible to get a response using TCP packets instead. TCP is the protocol used for all web browsers in addition to FTP, Telnet and others. This requires Administrative user rights, and on most operating systems also requires a helper library.

Timeout Speed

This option allows you to fine-tune your performance a little. When PingPlotter sends out a packet, it waits a certain amount of time for a response. The longer it waits, the more resources it needs to use (to keep sockets open), but the more likely that it will get a response. By default, PingPlotter will wait for 3 seconds for any packet to return. If the packet doesn't return in 3 seconds, then it is counted as a lost packet. If patience isn't one of your virtues, you can turn this down somewhat. No matter what your value is here, timed out packet will show with the time "9999."

Because of the performance enhancements offered by PingPlotter, it's unlikely that this option needs to be changed. If it's set too low, it can cause misleading data to be generated.

Time interval between hop traces

This can be an interesting number to manipulate. It's really meant for "advanced" users, so you don't NEED to change it.

PingPlotter sends out multiple packets at the same time and times everything at once. Actually, it leaves a tiny interval between each packet so as not to completely saturate your bandwidth when it sends out 30 packets. This time interval is adjusted by this parameter. Most of the time, 25ms is good. This falls within realm of what a 28.8 modem can perform. If you've adjusted your packet size, or your connection to the internet is really slow, you might want to crank this number up a little. If you have just oodles of bandwidth, you can crank it down a little. Be aware that a too-small number can adversely affect your data.

Packet Size

The Packet Size can make a considerable difference in latency performance. Normally, you want to use a relatively small number here. The default is 56 bytes, but in some cases you might need to lower this (especially on TCP port 80 packets, which sometimes get dropped unless they are 40 bytes). 1500 is lot of data, and should be used with great care. A 1500 byte packet means PingPlotter will be sending out 30-50 K per second worth of data, which can cause its own problems (and makes measuring latency more challenging).

ToS/DSCP header byte

This is relatively advanced course material – ie: if you don't know what this means, just leave it at 0.

The DSCP byte is often used by network providers to determine packet priority and sometimes make other decisions about the data. There may be an occasion where you want to manipulate this byte and test network performance. VoIP data, for example, is often characterized by a value in the DSCP byte.

This byte can be edited as Decimal, Hex, or Binary by hitting the button beside the edit field. The display on the button shows the current format this value is displayed in.

Note that under Windows XP and newer operating systems, a registry setting needs to be modified to allow this byte to be used. This is a system-wide setting that allows applications to write their own values into this byte. If your system needs to have this value set, PingPlotter will prompt to see if you want this setting changed (and a reboot will be required) as soon as you modify the value in this field.

Starting Hop

Sometimes hop 1 or 2 might never respond. Rather than continuing to pound away at these hops and never getting a response, it sometimes makes sense to just ignore the first hop or 2. This is totally normal on lots of cable modems, and can happen on any connection – where the first hop is always "silent". Ignoring the first non-responding hops will save some resources.

TCP Specific Settings

When using TCP packets, you can specify which target port to use. Usually, you'll want to use port 80 here, but you're welcome to use any reasonable port. Windows firewall blocks creation of TCP packets, so you'll need to use WinPCap to create packets under that OS (and possibly others).

Use non-threaded Name Lookups

This makes PingPlotter do all DNS lookups in the main application thread. If you're having DNS lookup problems, it might be worth trying. Some older Windows OS versions (Windows NT) may not like multiple name lookups being done. This also turns on a separate lookup if you're doing single-threaded tracing.

Maximum concurrent requests

PingPlotter sends multiple requests simultaneously, but we need to put a limit on it so we're not queuing more packets than we can send. Set this to 1 to only have one request out there at once. For ICMP types, this is concurrent **threads**. For UDP and TCP (which share a single thread for all packets), this is requests.

Setting this too high can cause PingPlotter to crash, especially on old operating systems. 45 is proven safe. If you have a really fast trace interval, you may need to increase this number to support the trace interval. Usually, you want to lower the "Timeout Speed" setting before raising this, as that might be adequate.

Cargo

Note: This is an advanced option that should probably not be changed. This is used to diagnose network problems when specific **data** is sent - which is a highly unlikely problem for most networks.

By default, PingPlotter (basically) pads the outgoing packet with repeating string representing the current PingPlotter version and name. An example of this is PingPlotter250, repeated over and over to fill the cargo. If you suspect that your network may be having problems when you send specific byte codes, you can enter the hex code that you want repeated, OR a link to a file to read the byte string from. The cargo space for the packet will be padded with this data. Use this in conjunction with the packet size to create the network scenario you're looking to duplicate.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.1.6 Auto-Save Options

The "Auto-Save" settings allow you to automatically save data or graphical images at an interval of your choice.

Auto-save data

Turning this setting on will cause PingPlotter to save a .pp2 file at the interval you choose. You should specify the directory you'd like to save in, and you should also include [some variable](#)⁵¹ as part of the name – otherwise every save will overwrite the previous save. This can be coupled with "Maximum samples to hold in memory" to minimize the amount of memory used on long monitoring projects. The "Save Interval" must have elapsed before the first save happens (this is so short-term tests don't auto-save data). A suggested number for this is 30 minutes.

The file name defaults to the current path. [Visit the topic on Variable Substitution for a list of variables](#)⁵¹ that can be used.

Make sure your file name will resolve to a valid file name. Remember that a colon (:) should only be used with a drive letter, and the / isn't a valid character.

A good example of a file name is: \$host \$date \$hour00

If the file already exists, it will be overwritten. This means you can combine a pretty fast save interval, and couple that with a longer time file name (ie: \$host) and save your data often without creating a lot of extra files.

Also, if you specify a directory that doesn't exist, then PingPlotter will try to create that directory (this is

useful if you want to save data about a specific host in a directory that matches that host name - you can use "c:\ppdata\%host%\%host %date" (or something similar) for your filename.

Auto-save image

PingPlotter can be set to automatically create an image of the current data (as seen on the graph). Using this, one could auto-save for a time period, then create some kind of "album" to post for proof of a problem. The "Save Interval" must have elapsed before the first save happens (this is so short-term tests don't automatically create images).

Note that the image that is saved is the same as the image on the PingPlotter display at the time the period elapses. If you've changed the focus of the PingPlotter display, then the current image saved rather than the current data at the time of save.

[See the Variable Substitution section](#) ^[51] for file name variables and tips.

Maximum samples held in memory.

This is the maximum number of samples that will be held in memory at any one time. Older samples are purged from memory. This can be coupled with auto-saving of data to keep your memory images small, but still have access to all the data collected.

PingPlotter averages about 44 bytes per sample (this can be more or less, but this is the memory for a 20 hop route), so 20,000 samples is still less than 1 megabyte of memory.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.1.7 Route Change Options

Default Settings: Route Masking

Exclusion Masks

191.168.1.255	Add IP Mask
	Expand selected mask
	Delete Mask

Show route change indicator when list hidden

List the IP "exclusions" for which you want to ignore route changes. If multiple address fit into the same mask, they will be treated as equivalent (and data will be treated as equivalent during data collection).

Use the format IP/bits (like 192.168.1.1/24) or IP/mask (like 192.168.1.1/255.255.255.0). Masks can be in IPv6 or IPv4 format, and use standard internet masking formats.

Example: if you have a normally oscillating route that oscillates between 192.168.1.1 and 192.168.2.1, then enter 192.168.1.1/255.255.253.255 in the exclusion list above - bit 2 of the third byte gets ignored.

If you want to ignore ALL route changes, enter the word ALL.

By default, PingPlotter tracks all route changes. Sometimes, these changes might be normal and might add confusion (ie: too much data) rather than clarity. PingPlotter supports "masking", or excluding certain changes.

Route Change Exclusion Masks

Route change masks can be manually added. Alternately (and much easier) is to add a mask by right-clicking on the oscillating router that you'd like to stop signaling a route change.

By default, any route change is recorded and noted by PingPlotter. In some cases, this may not be desired behavior - an example of this is if something in your regular trace route oscillates between 2 (or more) routers based on load.

If you're seeing route oscillation (where a specific hop regularly changes between 2 or 3 different IP addresses, but the rest of the route doesn't change), then you can add a mask to this list to suppress route change notifications when this happens.

To do this, hit "Add IP Mask" and enter the first IP address. Once it's in the list, select it, and then hit the next button down (Or with XX.XXX.XXX.XXX). In the popup, enter the next IP address of the oscillating set. Repeat this if there's a third one (or more).

Show route change indicator when list hidden

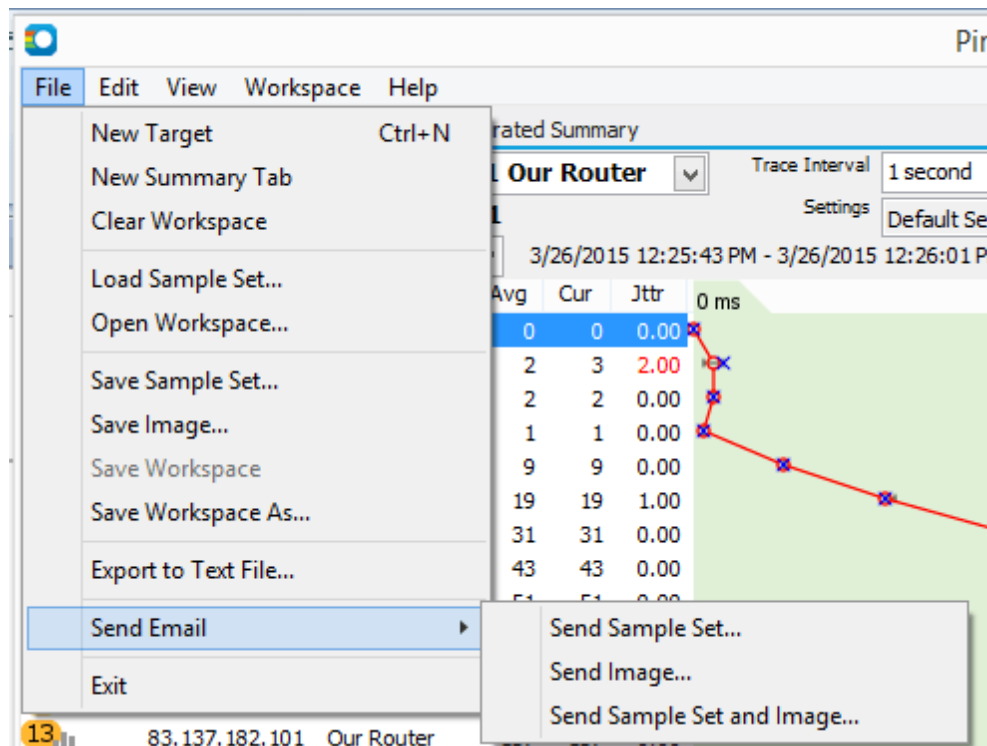
Normally, the route change window is hidden on the main screen. If this is the case, then the "Route Change" button will be highlighted whenever there is a route change - which you can click on to show the route change window.

If you're getting a lot of route thrashing that you don't want to know about, you can turn off this indicator here. This is *on* by default.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.2 Menus

8.2.1 File Menu

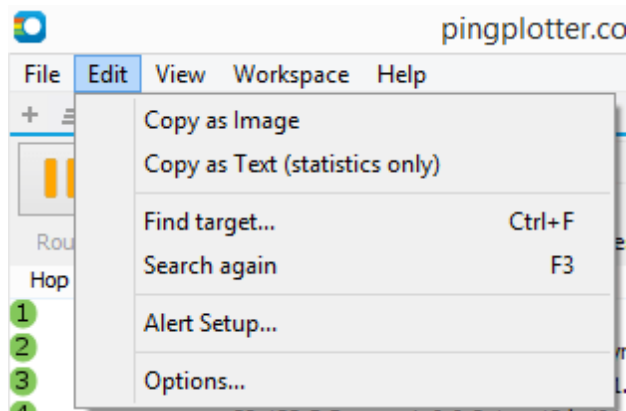


1. **New Target...** - This will create a new empty target area where you can trace to a new instance. See the documentation on "Tracing to Multiple Targets" for more details.
2. **New Summary Tab** - In PingPlotter Pro - this will create a new summary tab (we cover custom summary tabs in more detail here).
3. **Clear Workspace** - In PingPlotter pro - this will clear out your current workspace
4. **Load Sample Set...** - Loads a previously saved sample set. The default extension for PingPlotter saved sample files is .pp2, or PingPlotter save file format.

5. **Save Sample set...** - Allows you to save the current sample set to an external file. These files are saved in .pp2, or PingPlotter's save file format.
6. **Save Image..** - Saves the current graph in .png, .gif or .bmp format. See the [Autosave section](#)^[119] for information on how to automate the saving of graph images.
7. **Save Workspace/Save Workspace As...** - In PingPlotter Pro - this will save your current workspace (we cover workspaces in more detail here).
8. **Export to Text file...** - Exports trace data to a comma delimited text file. Click [here](#)^[51] for an explanation of the export options available from PingPlotter..
9. **Send Email** - This option will launch your email client and automatically create an email with the current sample set, image or both.
10. **Exit** - Exits PingPlotter. By default you'll be prompted to save your current sample set if you haven't done so already (click [here](#)^[107] to see how to change this option).

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.2.2 Edit Menu



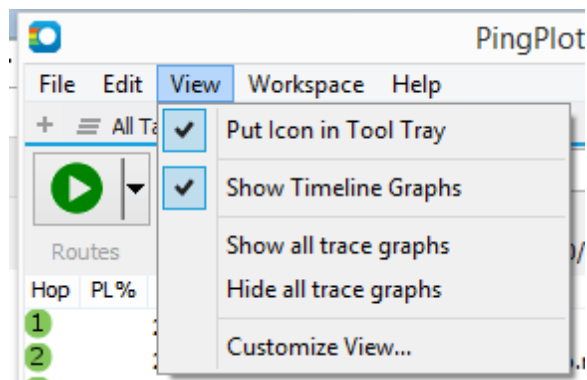
1. **Copy as Image** - Copy the current graph to the clipboard as an image. From here, you can paste the image into your favorite graphics program or an email.
2. **Copy as Text** - Copy the current graph to the clipboard as text. Hold down the shift key when clicking the Edit menu to toggle between copying all the collected data details, or copying a summary.
3. **Find target...** - In PingPlotter Pro - you can use this feature to search through your target list for a

specific host

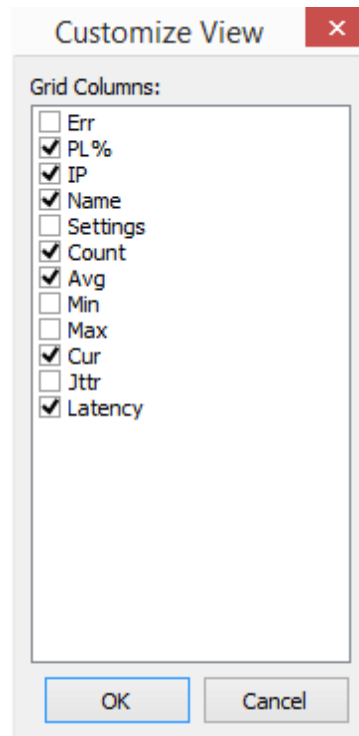
4. **Search again** - In PingPlotter Pro - this will perform another search using your last used search term
5. **Alert Setup...** - Create and Edit alert configurations. See the [Alert setup options](#)^[33] for more details.
6. **Options...** - Go to the configuration and options setup area.

Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

8.2.3 View Menu



1. **Put Icon in Tool Tray** - When this option is on, PingPlotter will move to the tool tray when minimized.
2. **Show Timeline Graphs** - Don't show/show the Timeline Graphs (the bottom history graphs).
3. **Show/Hide all trace graphs** - In PingPlotter Pro - this feature will open (or hide) all of your targets' trace graphs.
4. **Customize View...** - this will allow you to enable / disable columns from the upper trace graph.



Some of the features listed in this topic are only available in PingPlotter Pro and/or PingPlotter Standard. See our [product comparison](#) page for more details

Part



More Information

9 More Information

9.1 Support

We offer a number of great support resources for PingPlotter that might be able to answer your questions without having to ask us. We're always happy to get an email from you as well though.

Support Page:

<http://www.pingplotter.com/support.html>

Troubleshoot and Solve Your Network Problem:

<http://www.pingplotter.com/netnirvana/>

A Lineup of Common Network Problems:

<http://www.pingplotter.com/commonnetworkproblems/>

Getting Started Guide:

<http://www.pingplotter.com/gsg/>

Tutorial and Product Manual:

<http://www.pingplotter.com/manual>

Knowledge Base:

<http://www.pingman.com/kb/>

Support forums:

<http://www.pingman.com/forums>

Email Support:

support@pingplotter.com

9.2 Purchasing

PingPlotter has three different editions - Free, Standard, and Pro. To see more info about each version, be sure to check out our [feature comparison page](#).

PingPlotter Free doesn't require any sort of license key - and you're welcome to run that version of the program for as long as you'd like at no cost.

PingPlotter Standard and Pro both require a purchased license key to run after their 30-day evaluation period has expired. If no license is entered after the evaluation period has expired, the program will revert to PingPlotter Free.

We offer a variety of ways to purchase licenses, including credit card, PayPal, check, money order, or via purchase order (after approval). The very easiest way to order is with a credit card online – you'll get a license key immediately, and you won't have to wait for your order to be processed by hand (as some other methods require).

To purchase a license, visit our online order page at:

<http://www.pingplotter.com/order.html>

Thank you!

Index

- A -

Alerts Setup 33
Alerts Setup: Event Notification 44
Alerts Setup:Event - Launch an executable 47
Alerts Setup:Event - Log to file 47
Alerts Setup:Event - Play a Sound 44
Alerts Setup:Event - Send Email 45
Alerts Setup:Event - Tray Icon Change 48
Alerts Setup:Troubleshooting Alerts 41
Applying changes 106
Auto-save 53, 119

- C -

Cargo 115
Clipboard 25
Colors 112
Column Widths 17
Command Line Arguments 99

- D -

DNS Name 16
DSCP 115

- E -

Examples 83
Exporting 25, 51

- F -

Flash Drive 103

- H -

History Graph 17

- I -

Image creation 119
Introduction 2

- J -

Jitter 83
Jitter Display 112

- L -

Latency 12
License Key Entry 101
Licensing 127
Long term monitoring 53

- M -

Masking 121
Memory footprint 53, 119
Menu
 Edit 123
 File 122
 View 124
Minimize instead of close 107

- O -

Options 106, 112

- P -

Packet size 115
Packet type 115
Ping Thresholds 112
Proxy server 110

- R -

Registration 127
Removeable Drive 103
Reporting 25, 51
Reset and Restart 16

Round Trip Time 107
Route change masks 121
Route Changes 16

- S -

Samples to Include 16
Save on exit 107
Scrolling 21
Shortcut Keys 21
Show on summary graph 107
SMTP Password 111
SMTP Server 111
SMTP Username 111
Support 127

- T -

Time Graph 17
Time Graph Interval 100
Timeline Graph 17
Timeline Graphing 21
Timeout speed 115
Tool Tray 107
ToS 115
Trace Address 17
Trace Graph 17
Trace Interval 16
Trace Route 12
TTL 12

- U -

USB Drive 103

- V -

Version check 110
VoIP Troubleshooting 83

- W -

Whois 102
Whois lookup servers 110

