

Network Troubleshooting

HOW TO RUN A PING TEST IP

To run a ping test, we'll open the same command line and type `ping (host)` with *(host)* being the website you're trying to connect to. You'll get something like this (we'll use veesp.com again as an example):

ping vesp.com

Pinging vesp.com [104.22.74.140] with 32 bytes of data:

```
Reply from 104.22.74.140: bytes=32 time=ms TTL=56
Reply from 104.22.74.140: bytes=32 time=ms TTL=56
Reply from 104.22.74.140: bytes=32 time=ms TTL=56
Reply from 104.22.74.140: bytes=32 time=ms TTL=56
```

Ping statistics for 104.22.74.140:

```
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 1ms, Maximum = 2ms, Average = 1ms
```

As expected, Vesp is online and running. If you're having problems it could be on a hop, which is when you'd use the traceroute utility.

HOW TO USE THE MTR COMMAND

MTR stands for "my traceroute" and displays the route that a connection to a particular system takes by providing a continuously updating display of timing data.

Because MTR provides an image of the route the traffic takes from one host to another, it is essentially a directional tool. The route built between two points on the Internet may vary greatly under different circumstance like the location and the routers that are located upstream. For this reason, it is a good idea to collect MTR reports in both directions for all the hosts that are experiencing connectivity issues.

Vesp Customer Support will often request MTR reports both to and from your server if you are experiencing networking issues. This is because, from time to time, MTR reports will not detect errors from one direction when there is still packet loss from the opposite direction.

Install MTR Debian/Ubuntu

```
sudo apt update
sudo apt install mtr
```

Install MTR CentOS/RHEL/Fedora

```
yum install mtr
```

Install MTR Windows

You can download WinMTR here: <https://sourceforge.net/projects/winmtr/files/latest/download>

Generate **MTR** report using the following syntax:

mtr -rwn4 veesp.com

```
[root@localhost ~]# mtr -rwn4 veesp.com
```

HOST:	localhost.localdomain	Loss%	Snt	Last	Avg	Best	Wrst	StDev
1. --	212.6.44.1	0.0%	10	0.7	0.7	0.7	0.8	0.1
2. --	91.203.69.80	0.0%	10	0.6	0.9	0.6	3.3	0.8
3. --	91.203.69.73	0.0%	10	1.3	1.2	1.0	1.5	0.1
4. --	78.28.193.98	0.0%	10	1.4	1.5	1.3	1.6	0.1
5. --	213.248.84.32	0.0%	10	0.9	1.0	0.9	1.1	0.1
6. --	62.115.119.193	0.0%	10	2.2	1.8	1.3	3.9	0.8
7. --	195.12.254.187	0.0%	10	2.0	1.9	1.5	2.4	0.3
8. --	104.22.74.140	0.0%	10	1.0	1.0	0.8	1.1	0.1

Beyond simply seeing the path between servers that packets take to reach their host, MTR provides valuable statistics regarding the durability of that connection in the seven columns that follow.

Loss% column shows the percentage of packet loss at each hop.

Snt column counts the number of packets sent.

The next four columns **Last**, **Avg**, **Best**, and **Wrst** are all measurements of latency in milliseconds (e.g. ms). **Last** is the latency of the last packet sent, **Avg** is the average latency of all packets, while **Best** and **Wrst** display the best (shortest) and worst (longest) round trip time for a packet to this host. In most cases, the average (**Avg**) column should be the focus of your attention.

The final column, **StDev**, provides the standard deviation of the latencies to each host. The higher the standard deviation, the greater the difference is between measurements of latency. Standard deviation allows you to assess if the mean (average) provided represents the true center of the data set, or has been skewed by some sort of phenomena or measurement error. If the standard deviation is high, the latency measurements were inconsistent. After averaging the latencies of the 10 packets sent, the average looks normal but may in fact not represent the data very well. If the standard deviation is high, take a look at the best and worst latency measurements to make sure the average is a good representation of the actual latency and not the result of too much fluctuation.



The **r** option flag generates the report (short for --report).

The **w** option flag uses the long-version of the hostname so our technicians and you can see the full hostname of each hop (short for --report-wide)

The **n** option flag used for toggle DNS on/off. Each numbered line in the report represents a hop. Hops are the Internet nodes that packets pass through to get to their destination. The names for the hosts (e.g. "inner-cake" and "outer-cake" in the example) are determined by reverse DNS lookups

-4 | -6 Use IPv4 or IPv6 only