ping, traceroute and mtr Network path testing

START HERE

`ping -c 4 <host>`
Is it reachable?

`traceroute -n <host>`
Path, no DNS

`ping -4 <host>`
Force IPv4

`mtr <host>`
Live path and loss

`ping -6 <host>`
Force IPv6

`mtr -rwc 100 <host>`
Report to share

`traceroute <host>`
Show the path

`mtr -T -P 443 <host>`
Trace with TCP

TTL 1 router 1 replies → TTL+1 → **TTL 2** router 2 replies → TTL+1 → **TTL n** last router → reaches → **Target** final reply

Probes with rising TTL; each router answers with ICMP Time Exceeded

CONCEPTS

ping tests the round trip to one host; traceroute and mtr find where along the path it breaks or slows.

Many routers rate-limit or drop probe replies, so loss on one hop that does not continue to later hops is usually harmless.

traceroute defaults to UDP probes and may need root for some modes; mtr and ping use ICMP unless told otherwise.

PING BASICS

`ping <host>`
Ping until interrupted (Ctrl+C)

`ping -c <count> <host>`
Stop after <count> replies

`ping -4 <host>`
Use IPv4 only

`ping -6 <host>`
Use IPv6 only

`ping -i <seconds> <host>`
Wait between packets
Decimals allowed; very small values need root.

`ping -s <bytes> <host>`
Set payload size (default 56)

`ping -W <seconds> <host>`
Wait this long for each reply

`ping -w <seconds> <host>`
Give up after a total deadline
Exits after the deadline regardless of replies; handy in scripts.

`ping -n <host>`
Skip reverse DNS lookups

`ping -q -c <count> <host>`
Show only the summary

`ping -t <ttl> <host>`
Set the IP time to live

`ping -I <interface> <host>`
Send from an interface or address

`ping -D <host>`
Prefix each line with a timestamp

`ping -M do -s <bytes> <host>`
Probe path MTU, never fragment
Add 28 bytes (IP + ICMP headers) to the payload to get the packet size; Message too long means the MTU is smaller.

READING PING

Output fields

Field	Means
`icmp_seq`	Packet number; gaps mean loss
`ttl`	Remaining hops on the reply
`time`	Round-trip time in ms
`% packet loss`	Share of requests unanswered
`rtt min/avg/max/mdev`	Latency summary and jitter

Exit status

Code	Means
`0`	At least one reply
`1`	No reply received
`2`	Other error

A reply time that jumps on the first packet only is usually ARP or DNS warm-up, not a slow link.

TRACEROUTE

`traceroute <host>`
Trace the path (UDP probes)

`traceroute -n <host>`
Do not resolve names

`traceroute -I <host>`
Probe with ICMP echo
Use when UDP probes are filtered; may need root.

`traceroute -T -p <port> <host>`
Probe with TCP SYN to a port
Default port is 80; useful through firewalls that allow web traffic.

`traceroute -4 <host>`
Use IPv4 only

`traceroute -6 <host>`
Use IPv6 only

`traceroute -m <hops> <host>`
Set maximum hops (default 30)

`traceroute -f <ttl> <host>`
Start at this hop

`traceroute -q <count> <host>`
Probes per hop (default 3)

`traceroute -w <seconds> <host>`
Wait time for each reply

`traceroute -i <interface> <host>`
Send via an interface

TRACEROUTE (cont.)

`traceroute -s <address> <host>`
Use this source address

`traceroute -A <host>`
Show AS numbers per hop

Reading hop lines

Shown	Means
`* * *`	No answer; filtered or rate-limited
`three times`	RTT of each of the 3 probes
`!H / !N / !X`	Host, network or admin unreachable
`different IPs on a hop`	Load-balanced paths

MTR LIVE VIEW

`mtr <host>`
Live trace with loss and latency

`mtr -n <host>`
Show IPs only, no DNS

`mtr -b <host>`
Show names and IPs

`mtr -4 <host>`
Use IPv4 only

`mtr -6 <host>`
Use IPv6 only

`mtr -i <seconds> <host>`
Seconds between probes (default 1)

`mtr -s <bytes> <host>`
Probe packet size

`mtr -m <hops> <host>`
Maximum hops to probe

`mtr -f <ttl> <host>`
Start at this hop

`mtr -I <interface> <host>`
Send via an interface

`mtr -a <address> <host>`
Bind to a source address

`mtr -z <host>`
Show AS number per hop

Interactive keys

Key	Does
`d`	Cycle display mode
`n`	Toggle DNS
`p / Space`	Pause / resume
`r`	Reset counters
`j`	Toggle latency / jitter
`o`	Choose columns
`u`	Switch ICMP / UDP
`q`	Quit

MTR REPORT AND PROTOCOLS

`mtr -r -c <count> <host>`
Report mode: run, print, exit

`mtr -rw -c <count> <host>`
Wide report, full hostnames

`mtr -rn -c <count> <host>`
Numeric report

`mtr -T -P <port> <host>`
TCP SYN trace to a port
Packet size is ignored with TCP.

`mtr -u <host>`
UDP probes instead of ICMP

`mtr -o "LSD NBAW X" <host>`
Choose and order columns

`mtr -j -c <count> <host>`
JSON output

`mtr -C -c <count> <host>`
CSV output

Report columns

Column	Means
`Loss%`	Share of probes lost at this hop
`Snt`	Probes sent
`Last`	Latest RTT in ms
`Avg`	Mean RTT
`Best / Wrst`	Fastest / slowest RTT
`StDev`	Spread of RTT; high means jitter

Judge loss by the last hops: loss that appears at one hop and disappears further on is rate limiting; loss that continues to the target is real.

Run at least 50-100 cycles for a report, and run it from both ends when you can.

RECIPES

Is the network or DNS broken?

```
ping -c 3 1.1.1.1
ping -c 3 example.com
# IP works, name fails: DNS problem (see dig)
```

Evidence for your ISP

```
mtr -rwn -c 100 example.com > mtr.txt
mtr -rwn -T -P 443 -c 100 example.com >> mtr.txt
```

Find the path MTU

```
ping -c 2 -M do -s 1472 <host>
# fails? lower -s until it passes; MTU = size + 28
```

Watch for drops over time

```
ping -D -i 0.5 <host> | tee ping.log
grep -c 'bytes from' ping.log
```