

Name

tulip – a tool for user-level internet path diagnosis

Synopsis

tulip <reordering|loss|queuing> [options] <destination>

Description

Tulip is a tool to localize the occurrence of packet reordering, loss and queuing delay along the path from the source to an arbitrary destination. It leverages ICMP timestamps and sequential IP identifiers, two common features in today's routers.

Tulip is implemented on top of Scriptroute (<http://www.scriptroute.org>), which needs to be installed (don't need to enable remote execution) before tulip can be used. Invoke tulip with the path to srinterpreter if srinterpreter is not installed in /usr/bin.

Options**--count COUNT**

Number of measurements to use with each router on the path.

--lag LAG

Milliseconds to wait between the start of successive measurement rounds. A round consists of one measurement to each router on the path. Only one router is measured at a time to control network load. Scriptroute rate-limiting or high round trip times to routers may mean that a round cannot be completed in LAG milliseconds. Use

--spread SPREAD

Millisecond separation between probes in a single measurement. The default is zero (back-to-back) probes.

--start START

Hop number to start diagnosis from. The default is 1.

--end END

Hop number to end diagnosis at. The default is the last hop in the traceroute.

--noskip

Stop querying rate-limiting routers. The default is skipping. Rate-limiting routers are identified as those with too many non-returned responses.

--printfrequency FREQUENCY

Number of measurements after which to print the results. The default is COUNT.

--noallhops

Use round trip measurements for hops that don't support forward primitives. The default is to probe all hops.

--dstonly

Probe end-to-end path only. By default, routers on the path are also probed.

--nonames

Do not resolve ip addresses. By default, the IP addresses are resolved.

--prefixpath PP

Number of initial hops to excuse for local load balancing. The default is 0.

--help Display a quick help message.**--verbose VERBOSE**

Level of verbosity (0-10). The default is 0 (minimum).

Output Format

```
% tulip reordering --count 5 www.iitd.ernet.in
// path to the destination
---- after 5 measurements ----
1 eureka-GE1-5.cac.washington.edu (128.208.4.100) rt=0.000 (0/5) fw=0.000 (0/5)
2 uwbr2-GE0-1.cac.washington.edu (140.142.150.24) rt=0.000 (0/5)
//remaining hops
%
rt is the round trip reordering rate and fw is the forward path reordering rate. No forwarding path rate is
printed when the router does not support sequential IP-IDs.
```

```
% tulip loss --count 100 www.iitd.ernet.in
//traced path
---- after 100 measurements ----
1 eureka-GE1-5.cac.washington.edu (128.208.4.100) rt=0.000 (0/100) fw=0.000 (0/100) co=0.950 (95/100)
ro=0.000 (0/100)
2 uwbr2-GE2-0.cac.washington.edu (140.142.153.24) rt=0.000 (0/100)
//remaining hops
%
rt denotes the round trip path to the remote end. Round trip loss implies that the data (middle) probe was
lost without losing any of the control probes. fw denotes the forward path to the router -- it counts the num-
bers of measurements in which the data probe was lost and the control probes are returned with consecutive
IP-IDs (the denominator in the computation is the number of measurements for which none of the probes
were lost), co is the number of measurements in which the IP-IDs were consecutive -- consecutive IP-IDs
are needed for forward path loss measurement. ro counts the number of measurements in which the probes
were reordered -- reordered probes can lead to incorrect inferences; increase the spread of the probes to
reduce reordering.
```

```
% tulip queuing --count 100 www.google.com
---- after 100 measurements ----
1 eureka-GE1-5.cac.washington.edu (128.208.4.100) rt: median=0.026 min=0.130 max=76.292
2 uwbr2-GE0-1.cac.washington.edu (140.142.150.24) rt: median=0.044 min=0.340 max=41.931
3 cnsp2-wes-ge0-1-0-0.pnw-gigapop.net (198.107.150.7) rt: median=0.110 min=0.610 max=10.875 fw:
median=0.000 max=11 quality=1.000 (0/0)
//remaining hops
%
rt is the round trip path: median is the median delay above the minimum, min is the measured minimum
round trip delay, and max is the measured maximum round trip delay. fw is the forward path to the router:
median is the median delay above the minimum delay (-1 implies that the clock calibration failed because
the router clock is not well-behaved or the path is too noisy; using more probes might help), max is the
maximum forward path delay above the minimum delay, quality is the quality of clock calibration. tulip
uses fixclock (written by Raphael Ryger as part of cing) for clock calibration; low quality may mean that
the clock calibration is suspect. In case you are curious, the two numbers in the parenthesis represent estimated
min round trip time based on clock calibration and the measured min round trip time.
```

Notes (bugs/todo)

Send bug reports and suggestions to Ratul Mahajan <ratul@cs.washington.edu>

Loss rate fraction is not intuitive.

Binary search is not implemented in this version.

Adapt fixclock's segmentsecs based on lag.