

How Traceroute Explains the Internet



TEAM CYMRU™


John Kristoff
jtk@cymru.com

History

- Van Jacobson releases traceroute.tar.Z

To: `ietf@venera.isi.edu,`
`end2end-interest@venera.isi.edu`
Subject: 4BSD routing diagnostic tool
available for ftp
Date: Tue, 20 Dec 88 05:13:28 PST
From: Van Jacobson

Traceroute Operational Review



UDP Traceroute

- Also known as UNIX traceroute
- ICMP TTL exceeded responses to ICMP messages?
- Probes initialize to an unlikely UDP destination port
 - Default begins at 33434
 - Incremented for each probe packet
- An ICMP port unreachable from the target is the goal

ICMP Traceroute

- Most commonly associates with Microsoft Windows
- Sends ICMP echo request probes
- An ICMP echo response from the target is the goal
- NOTE: ICMP TTL exceeded messages not a problem
- But filtering and probe response suppression can be

Anomaly: Multi-path

```
$ traceroute www.chinog.org
traceroute to www.chinog.org (74.208.62.118), 30 hops max,
                                         60 byte packets
```

...

```
4  te0-3-0-2.agr22.ord01.atlas.cogentco.com (154.24.4.41)      1.422 ms
   te0-3-0-2.agr21.ord01.atlas.cogentco.com (154.24.4.37)      1.175 ms
   te0-3-0-2.agr22.ord01.atlas.cogentco.com (154.24.4.41)      1.329 ms
5  be2524.ccr42.ord01.atlas.cogentco.com    (154.54.81.109)   1.475 ms
   be2521.ccr41.ord01.atlas.cogentco.com    (154.54.80.253)   1.710 ms
   be2522.ccr42.ord01.atlas.cogentco.com    (154.54.81.61)   1.455 ms
```

...

```
13  perfora.net (74.208.62.118)  17.273 ms  17.490 ms  17.276 ms
```

Anomaly: Missing Hop(s)

```
$ traceroute -n facebook.com
traceroute to facebook.com (173.252.120.6) ...

...

 7  31.13.25.32    34.084 ms
    31.13.25.106  34.137 ms
    31.13.27.40   33.957 ms
 8  204.15.23.247  33.548 ms
    31.13.27.133  33.785 ms
    173.252.64.65 33.694 ms
 9  * * *
10  * * *
11 173.252.120.6  33.786 ms  33.570 ms  33.614 ms
```

Anomaly: Unresponsive Target

```
$ traceroute -n -q1 www.northwestern.edu
traceroute to www.northwestern.edu
(129.105.215.254), 30 hops max, 60 byte packets
```

...

```
4  199.249.169.5    0.884 ms
5  129.105.247.224  1.418 ms
6  129.105.253.153  1.682 ms
7  129.105.247.97   2.006 ms
8  *
9  *
10 *
11 *
12 *
```


Hack: Uncovering a Target

```
$ sudo traceroute -T -n -q1 www.northwestern.edu
traceroute to www.northwestern.edu
(129.105.215.254), 30 hops max, 60 byte packets
```

...

```
4  199.249.169.5    0.797 ms
5  129.105.247.224  1.026 ms
6  129.105.253.153  1.662 ms
7  129.105.247.97   655.220 ms
8  129.105.46.196   1.896 ms
9  129.105.215.254  1.996 ms
```

Feature: MPLS option

```
traceroute -en -q1 www.google.com
```

```
traceroute to www.google.com (64.233.160.99), 30 hops max,  
60 byte packets
```

```
...
```

```
6 209.85.254.120 1.529 ms  
7 209.85.254.238 <MPLS:L=284331,E=4,S=1,T=1> 18.373 ms  
8 209.85.251.18 <MPLS:L=319427,E=4,S=1,T=1> 18.165 ms  
9 72.14.236.1 <MPLS:L=770010,E=4,S=1,T=1> 21.186 ms  
10 209.85.248.7 19.892 ms  
11 64.233.160.99 16.332 ms
```

Feature: Visualization (geoloc)

Open Visual Trace Route 1.5.0


Traceroute Sniffer Atheros L1C... 2D/3D google.com Timeout 10

Open Visual Trace Route 1.5.0 License

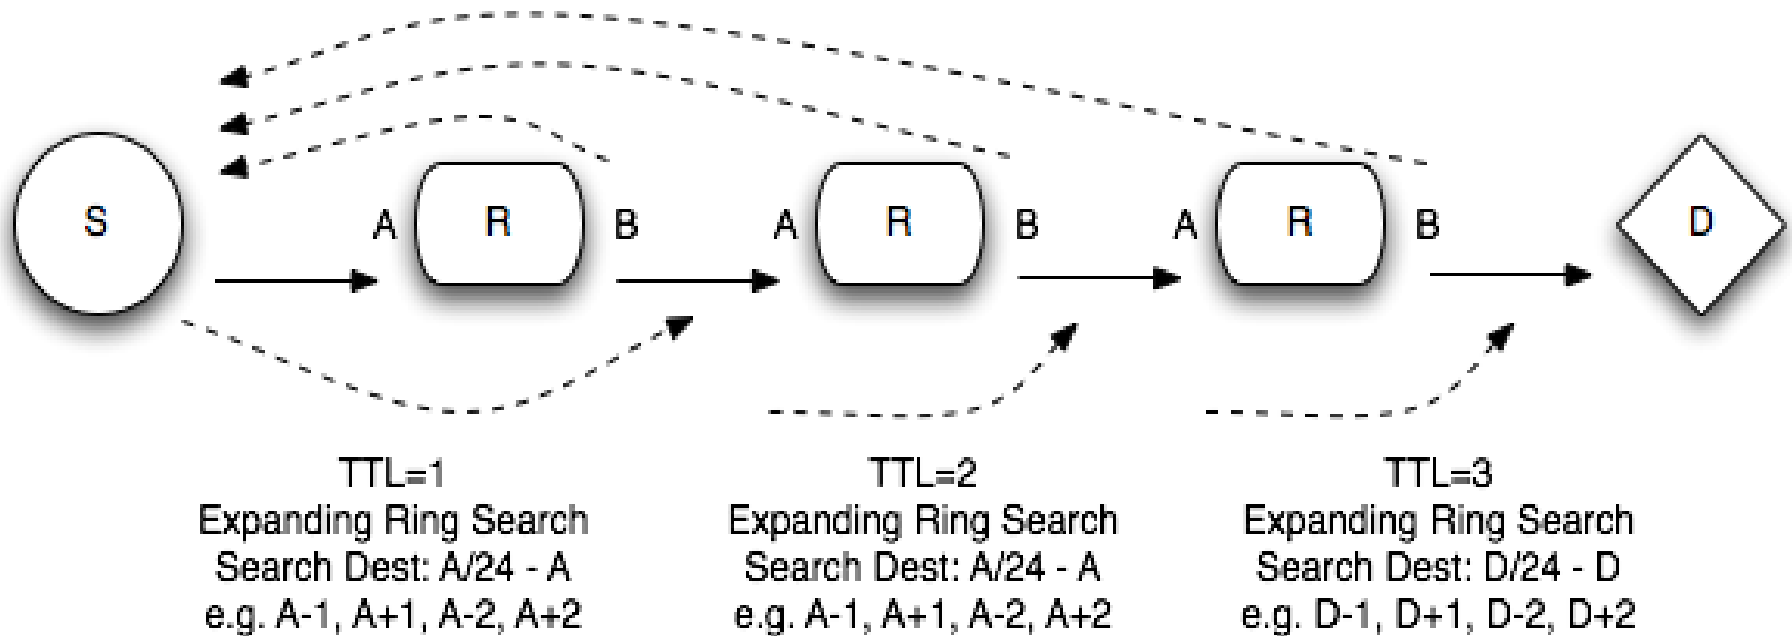
#	Country	Town	Lat	Lon	IP	Hostname	Lat...	Dns I...	Dista...
1	Japan	(Unknown)	36.0	138.0			<1	<1	0
2	Japan	Matsumoto	36.233307	137.9667			7	43	26
3	Japan	Matsumoto	36.233307	137.9667			<1	39	0
4	Japan	Kamakura	35.3089	139.5503			5	43	176
5	United States	The Dalles	45.544693	-121.1543	72.14.212.73	(None)	2	54	7939
6	Japan	(Unknown)	36.0	138.0	202.213.197.99	(None)	<1	82	7988
7	United States	Mountain View	37.419205	-122.0574	72.14.236.82	(None)	2	47	8431
8	United States	Mountain View	37.419205	-122.0574	209.85.244.67	(None)	<1	49	0
9	United States	Mountain View	37.419205	-122.0574	173.194.126.137	nrt04s05-in-f9.1e100.net.	3	48	0

Trace route completed in 0.501 s. Route length: 24,560 kms

Hack: Peer Discovery



Hack: Reverse Traceroute



Let's Take Stock

- Key figures, organizations, mailing lists, software
- Routers+routing, forwarding, and loop prevention
- ASNs, BGP peering policies, asymmetric paths
- UDP (TCP), source+destination ports, ICMP
- Domain name system, naming conventions
- Packet filtering and security policies
- Performance, round trip time, geo-location

...and much more

- IP headers and options
- ARP, MPLS, IPv4 versus IPv6
- An interface versus a host
- NAT
- Applications (ports)
- Router CPU protection, rate limiting
- ...and even the end-to-end argument

A bit of a stretch? Maybe not

- IP multicast
- DDoS
- Buffering, congestion avoidance and flow control
- CoS/QoS
- PKI – OK, maybe this one is a bit of a reach

Conclusion: Some References

- Richard A Steenbergen (RAS)
 - <http://cluepon.net/ras/traceroute.pdf>
 - A Practical Guide to (Correctly) Troubleshooting with Traceroute, ARIN on-the-road, Orlando 2015
- <http://kb.pert.geant.net/PERTKB/VanJacobsonTraceroute>
- Traceroute Anomalies, Martin Erich Jobst
- Traceroute Probe Method and Forward IP Path Inference, Lukie, Hyun, Huffaker
- Reverse Traceroute, Bassett, et al.