



OnTheInternet

THIS ARTICLE

DISCUSS

EMAIL TO A FRIEND

PRINTABLE VERSION

MORE INFORMATION

There are three commercial geographic traceroute applications:

[*GeoBoy by NDG Software](#)

[*NeoTrace by NeoWorx](#)

[*VisualRoute by Datametrics Systems Corporation](#)

Other notable traceroute utilities include:

[*GTrace by Cooperative Association for Internet Data Analysis \(CAIDA\)](#)

[*WhatRoute by Bryan Christianson.](#) This is available for the Mac.

[*Tracemap by Matrix and Directory Services \(MIDS\).](#) This is a Web-based geographic

Mapping Where the Data Flows

By Martin Dodge
m.dodge@ucl.ac.uk
<http://www.cybergeography.org>



When I type the Internet Society's Web address into my browser, the html and graphics are seamlessly downloaded and the page displayed within a couple of seconds. But how does this data actually get from the ISOC server to my PC in a basement office in central London? How does it flow through the Internet to reach me? You can answer this with traceroute, a useful tool that allows you to lift the lid on the Internet and get a packets'-eye view of the network. If you're like me—intrigued as to how the Internet works beyond your browser and the telephone jack in the wall—then traceroute can be a fun tool to use to explore and map the Internet.¹

A traceroute utility maps the path that data packets take between two points on the Internet, showing all of the intermediate nodes traversed, along with an indication of the speed of travel. Traceroute was invented in 1988 by Van Jacobson at Lawrence Berkeley National Laboratory in the U.S. Today a traceroute utility often comes as part of the operating system. Windows, for example, has a small utility called tracert, which is used by typing, at the MS-DOS prompt, tracert <Internet URL, e.g., www.yahoo.com>.

To illustrate how traceroute can map the Internet, I used it to analyze the path from my PC in London to www.isoc.org. The program works away for a few seconds as it dynamically explores the pathway the data takes, moving one node at a time. The end result is the following rather cryptic-looking output:

```
Tracing route to info.isoc.org [198.6.250.9]
over a maximum of 30 hops:
  1  10 ms  <10 ms  <10 ms  cisco-2.bart.ucl.ac.uk [128.40.59.245]
  2  <10 ms  <10 ms  <10 ms  128.40.100.254
  3  <10 ms  <10 ms  <10 ms  atmr-ulcc.lonman.net.uk [194.83.100.62]
  4  <10 ms  <10 ms  <10 ms  south-east-gw.ja.net [146.97.250.129]
  5  <10 ms  <10 ms  <10 ms  us-gw2.ja.net [128.86.1.242]
  6  80 ms   80 ms   80 ms  ny-pop.ja.net [193.62.157.10]
  7  80 ms   80 ms   80 ms  if-10-1.core1.NewYork.Teleglobe.net [207.45.196.157]
  8  70 ms   80 ms   80 ms  if-4-0-0.bb2.NewYork.Teleglobe.net [207.45.223.165]
  9  80 ms   80 ms   80 ms  Hsa12-0-0.GW1.NYC1.ALTER.NET [157.130.6.125]
 10  71 ms   80 ms   90 ms  159.ATM2-0.XR1.NYC4.ALTER.NET [146.188.180.34]
```

traceroute.

Xtraceroute by
Björn Augustsson

QUICK SEARCH

(Search disabled
for preview)



This might look meaningless to those unfamiliar with traceroute, but it is in fact a kind of one-dimensional map of how the data flows, with each node listed on a separate line. The map gives valuable information on the real-time routing of data packets through the Internet between London and Reston, Virginia, in the U.S.—the apparent location of the ISOC Web server. Each row of the chart contains a node number followed by three time measurements in milliseconds—such as 20ms 30ms <10ms. These are three separate measurements of the time it took the data packet to travel from the origin computer to that particular node and back. This is called the round-trip time (RTT) and gives an indication of the speed of each link. Finally, each row identifies a node as a domain name and numeric IP address.

The traceroute output is comparable to a railway timetable that charts (1) how trains travel through the network, (2) which stations they visit, and (3) when they arrive. The traceroute timetable above shows that data travelling from London to ISOC had to pass through 16 intermediate stations, or network routers, to reach the end of the line. To decode the traceroute chart further, read it line by line to see each section of the route through the network.

In our trace we begin at node one, which is the departmental router for my office, known by the somewhat cryptic domain name *cisco-2.bart.ucl.ac.uk*. The data packets move rapidly onward to the next node, which has only a numeric address and is likely to be an anonymous but important link somewhere in a server room in the university. At node number three the data leaves my university's internal network and joins the London metropolitan-area network—<http://www.lonman.net> [..uk/images/map99.gif](http://www.lonman.net/images/map99.gif)—which provides a fast backbone for universities and colleges in London. From there it moves on to the JANET backbone, the U.K.'s academic and research network—<http://www.ja.net/>—at a gateway machine called *south-east-gw.ja.net*. Node five is the gateway router to the transatlantic link for the JANET network, which connects to *ny-pop.janet*, its point of presence (pop) in New York. At this point the data packets have crossed the Atlantic, and consequently, you can see a marked jump in the round-trip travel times from under 10 milliseconds to 70 milliseconds caused largely by the 3,000-mile distance between London and New York. The traffic then flows into the Teleglobe backbone network—<http://www.teleglobe.com/>—in New York at nodes seven and eight before passing to AlterNet—part of UUNET's backbone empire, <http://www.alter.net/lang..en/network/>—also in New York.

Notice the strange, long domain names of these routers at the

core of the Internet. These names often hint at the city where the node is located. They could contain the full name of the city—such as *newyork*—or just an abbreviation such as *nyc*—which can require some educated guesswork to decode. Fortunately for traceroute explorers, many of the large backbone operators use similar naming conventions for their network infrastructures.

From AlterNet's network in New York, packets flow through two more nodes before dropping down to Washington, D.C., at node 12. Then it is on to two nodes in TCO, according to their domain name, which is most likely to be Tysons Corner, Virginia. We are nearing our target. At node 15 we have reached ISOC's network, and node 16 is the end of the line—the Internet Society's Web server. The ISOC home page took 16 hops, across four different networks, to get to my browser—quite a feat of routing and cooperation, but all in a day's work for the Internet. This happens, unseen, for the millions of Web surfers who need not worry about where the data flows.

Triangulating the Internet with Web-Based Traceroutes

Conventional traceroute utilities are limited in one important respect: the origin point of the exploration is fixed to the location of the PC running the trace. To overcome this limitation, you can use Web-based traceroutes that allow you to run a trace from different starting points. These allow you to explore the Internet's topology from multiple locations—a kind of virtual triangulation. I ran a traceroute from Canberra, Australia, to ISOC by using a Web-based traceroute publicly provided by Telstra, a major Australian telecom company and Internet backbone operator: <http://www.telstra.com.au/>. The output trace timetable is as follows:

```

traceroute to info.isoc.org (198.6.250.9), 30 hops max, 40 byte packets
 1 Ethernet0.dickson.Canberra.telstra.net (203.50.0.1) 1.78ms 2.32ms 1.60ms
 2 Serial6-0.civ2.Canberra.telstra.net (139.130.235.5) 2.51ms 3.04ms 2.24ms
 3 Fddi0-0.civ-core1.Canberra.telstra.net (139.130.235.226) 2.40ms 2.08ms 2.884ms
 4 Hssi0-1-0.pad-core3.Sydney.telstra.net (139.130.249.33) 7.73ms 7.801ms 8.827ms
 5 FastEthernet1-0-0.pad16.Sydney.telstra.net (139.130.249.235) 8.87ms 8.41ms 10.67ms
 6 205.174.74.33 (205.174.74.33) 220.48ms 218.85ms 218.83ms
 7 199.37.127.66 (199.37.127.66) 218.47ms 221.52ms 220.49ms
 8 POS2-0-0.GW5.SFO4.ALTER.NET (157.130.197.77) 219.04ms 219.79ms 239.09ms
 9 145.ATM2-0.XR1.SFO4.ALTER.NET (146.188.149.42) 221.84ms 219.39ms 220.68ms
10 291.at-5-0-0.TRJ.SCL1.ALTER.NET (152.63.48.210) 221.94ms 221.36ms 223.92ms
11 152.63.3.193 (152.63.3.193) 293.08ms 290.29ms 289.89ms
12 152.63.32.201 (152.63.32.201) 275.50ms 274.35ms 276.10ms
13 146.188.160.33 (146.188.160.33) 295.56ms 285.49ms 286.68ms
14 charlie.isoc.org (198.6.250.1) 293.60ms * 295.17ms
15 info.isoc.org (198.6.250.9) 294.96ms 293.06ms

```

Note that the output from this version of traceroute is in a slightly different format from the previous version.

The traceroute utility is installed on a Telstra server that is quite likely located in Canberra, given its domain name—*Canberra.telstra.net*—so this is where our data packets begin their journey to ISOC. The next two nodes in the trace are also within Canberra, according to their domain names. At node four the data moves a couple of hundred miles from Canberra

to Sydney.

The big trans-Pacific hop occurs at node six. There is certainly a marked jump in the RTT at this point in our journey, caused by the 7,500-mile distance across the Pacific Ocean. There is no domain name for node six to give us a clue to its location, but it is likely to be in California.

At node eight, the data joins the AboveNet network: <http://www.above.net/network/network.html>. The SFO domain name means it is probably in San Francisco. After another node in SFO, the data moves to a node in SCL, which is probably Santa Clara, California. The next major element in the journey is the hop across the continental U.S. from California to the Washington, D.C., region, at node 11. This long distance is matched by another significant increase in the RTT. Nodes 11–13 have no domains to tell us which networks they are on, but at node 14 the data has arrived at the ISOC network and Web server.

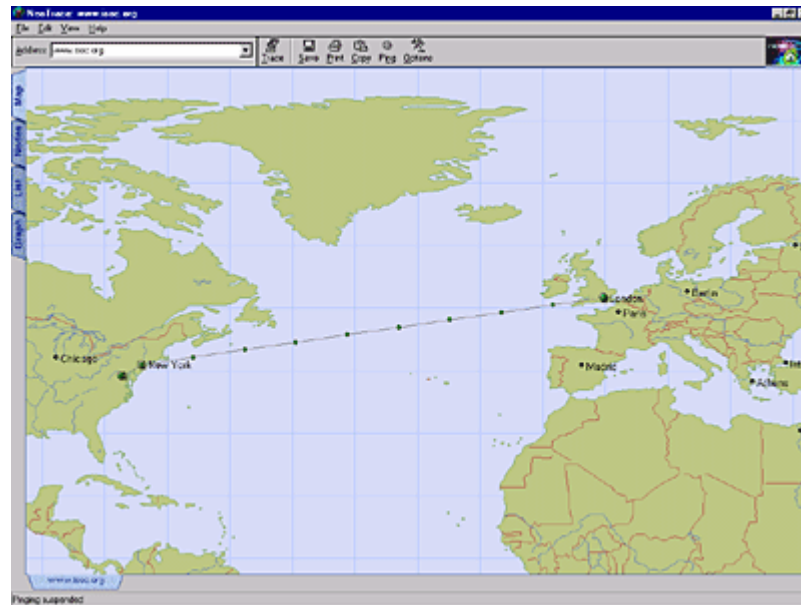
To run a trace from Canberra yourself, go to <http://www.telstra.net/cgi-bin/trace>. There are several hundred freely available Web traceroute servers in many different countries and cities across the world. Thomas Kernen maintains a good list of them at <http://www.traceroute.org/>.

Add a Bit of Geography

An obvious refinement of the basic traceroute is to show on a map the route the data takes visually. This is known as a geographical traceroute. A number of these applications (see sidebox) attempt, with varying degrees of success, to map the physical location of Internet nodes traversed in a trace. I used two of the best—NeoTrace and VisualRoute—to run the first trace example from London to ISOC to see how they performed when actually mapping the route. Both of them are easy-to-use, affordable application utilities.

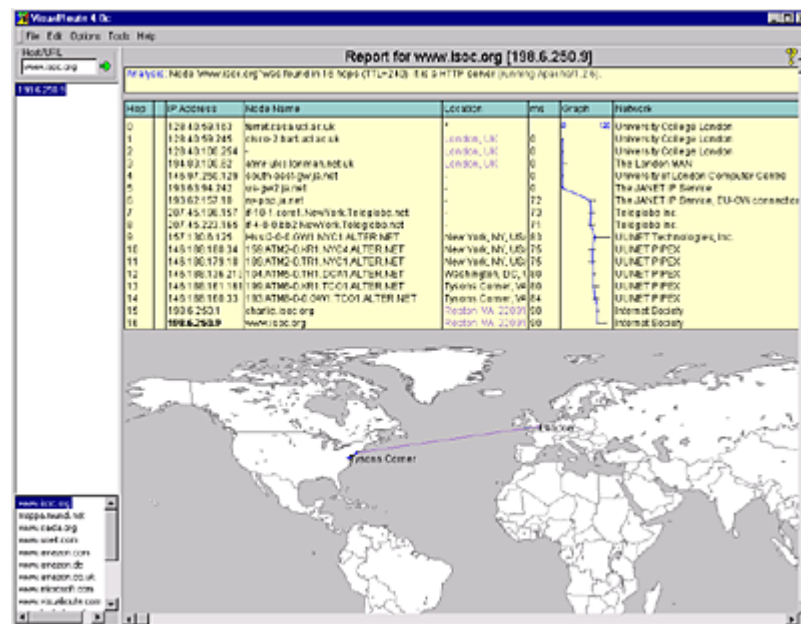
NeoTrace

NeoTrace is a geographical traceroute utility developed by NeoWorx. It provides four different views of the trace in tabbed panels: a geographic map, a nodal graph layout, a conventional listing, and a line graph of RTT performance. A trial of NeoTrace for Windows 95/98/NT can be downloaded from <http://www.neotrace.com> for free; the full program costs US\$ 29.95. The figure shows the map view result of our trace; it has done a good job, successfully locating and mapping 8 of 16 nodes, including the target in Reston.



VisualRoute

The second geographic traceroute application is Datametrics Systems Corporation's VisualRoute. The application interface includes both a zoomable world map and a detailed listing of the trace. A trial version of VisualRoute can be downloaded for free from <http://www.visualroute.com>, and the full program costs US\$ 29.95; it is available for Windows 95/98/NT, Linux and Solaris. VisualRoute also provides a Web server version, with examples currently located in England; the Netherlands; Canada; Freemont, California; and Fairfax, Virginia (<http://www.visualroute.com/server.html>). In our trace to ISOC, VisualRoute found and plotted the location of six nodes with confidence and made educated guesses at five others, including the target. The results are shown in the figure below.



Although geographic traceroutes are useful, their mapping

capabilities currently are limited by the serious difficulty in mapping an Internet node to an actual latitude and longitude. There is no automatic way to match up these virtual and real-world addresses. This is true even in the U.S., where one can at least match backbone nodes to particular cities somewhat reliably. Consequently, geographic traceroutes use a variety of heuristics to try to resolve a network node to a geographic location with various levels of accuracy.² This is a difficult problem to crack, and traceroutes largely depend on looking up Internet addresses in static databases of latitude and longitude. These databases, however, are only partial and cannot keep pace with the Internet's constant growth and change. They seem especially weak outside the U.S.

NeoTrace and VisualRoute offer a partial solution, allowing you to add your own nodes and locations to their databases, along with more detailed maps. A more effective solution is to add geographic location information to domain names, as set out in a DNS-LOC proposal that can be dynamically queried, but this has not been implemented widely.³ An alternative might be the NetGeo, developed by the Cooperative Association for Internet Data Analysis (CAIDA), a service for mapping network entities—such as IP addresses, domain names, and ASes—to geographical locations.⁴

Conclusions

Why is traceroute useful? First, traceroute is an important Internet debugging tool for those involved in keeping networks running. It can help identify routing problems quickly and simply. It can also be useful in tracking down the source of spam e-mail,⁵ as well as in trying to find a Web site's true location before giving it one's credit card details. Many Web sites using country-level domains are not actually hosted in the nation they indicate.

On another level, traceroute can help satisfy those who are curious to know how their computers connect to the Internet and how they can access information from all around the world as if it were just next door. Traceroute reveals the hidden complexity of data's path to a given destination—sometimes across 10 or 20 nodes or more, perhaps owned and operated by competing companies. Looking at what is happening in real time on the Internet always gives me a small sense of wonder that this system works so well, enabling tens of millions of people to communicate daily. So if you have a little time, go and trace the route to your favorite Web site and uncover the hidden complexity of the Internet that lies beyond the browser window. You may be surprised by where the data flows.

References

¹ The following articles provide good background information on using traceroute to explore the Internet:

J. Rickard, "Mapping the Internet with Traceroute," *Boardwatch* magazine, December 1996.

<http://boardwatch.internet.com/mag/96/dec/bwm38.html>.

J. Carl, "Nailing Down Your Backbone: The Imprecise Art of Tracerouting," *Boardwatch ISP Directory*, Summer 1999.

<http://boardwatch.internet.com/isp/summer99/tracerouting.html>.

S. Dumett, "Tracing Your Route Through The Net," *PreText* magazine, March 1998.

<http://www.pretext.com/mar98/columns/posttext.htm>.

Unix Man page for Traceroute. <http://www.zytek.com/traceroute.man.html>.

2 For a nice summary of various means of determining an Internet host's geographic location, see "Finding a Host's Geographical Location," by Uri Raz. <http://www.private.org.il/IP2geo.html>.

3 "Geo-enabling the Domain Name System."

<http://www.ckdhr.com/dns-loc/> ..

4 "NetGeo—the Internet Geographic Database." CAIDA.

<http://www.caida.org/Tools/NetGeo/>.

5 B. Mattocks, "Reading Email Headers."

<http://www.blighty.com/spam/bill.html>.

[TOP](#) | [HOME](#) | [NEWS](#) | [THEMES](#) | [ARCHIVES](#) | [ABOUT E-OTI](#) | [SEARCH](#)