Understanding Traceroute

One of the most effective tools for hunting down problems on an errant network is probably already installed on your machine. Traceroute can be useful to help solve problems for you and your users.

By Clive Grace

Bandwidth-hogging games aside, network outages open up a can of worms for network support, and knowing where (and how) to look for network problems is often the first and best way to form a solution.

One of the most commonly-used and freely available tools at your disposal is Traceroute, which uses TCP/IP to carry and return data relating to packets that describes how and where data originates and (in some cases) returns from.

Protocols

TCP/IP is the *lingua franca* of the Internet. It's the way the Internet enables processes on different host computers to communicate using a "connectionless" method through which data is delivered from machine to machine. Connectionless protocols send data down the line without first setting up a connection (as opposed to connection-oriented protocols which establish a connection before data is transmitted).

Below TCP/IP sit various methods (sometimes called media protocols) such as Ethernet that help shunt data along specific routes intended for specific destinations. Above TCP/IP lie protocols used by applications such as HTTP (HyperText Transport Protocol) for Web browsers, NNTP (Network News Transport Protocol) for Usenet news readers and SMTP (Simple Mail Transport Protocol) for the transport of Mail.

These protocols are managed by applications your users are using. Whether it's a Web browser, a server, an email client, or even generic TCP/IP support built into Windows 95 (such as application auto-updaters), the protocol exists so that each application group can exchange data with the corresponding servers using TCP/IP.

When users send email messages or request a connection to a Web server, the message is broken into packets and transmitted. These packets travel along different paths depending on their type, whereupon the complete message is rebuilt at its destination.

You can trace the route these packets take by using Traceroute.

```
Microsoft(R) Windows 98
   (C)Copyright Microsoft Corp 1981-1998.

C:\WINDOWS>tracert

Usage: tracert [-d] [-h maximum_hops] [-j host-list] [-w timeout] target_name

Options:
    -d                 Do not resolve addresses to hostnames.
    -h maximum_hops    Maximum number of hops to search for target.
    -j host-list       Loose source route along host-list.
    -w timeout         Wait timeout milliseconds for each reply.

C:\WINDOWS>
```

Figure 1 - Output from the Windows 9x Tracert command.

What Is Traceroute?

Traceroute started life as a simple command-line diagnostic tool for local area networks. Results are presented as rows of information, with each line showing a node the packet visited along the way to its destination. Each node represents one set of choices (or hops) made by successive routers to deliver data. Data might not always follow the same path, even to the same destination, and it is for this reason that Traceroute is invaluable in hunting down slow or non-existent network connections.

How It Works

Traceroute's function is very simple. Although some PC packages make very good use of the PC's graphical interfaces - literally mapping the connection paths between nodes - these implementations rarely vary dramatically from the Unix sources from which they were built (in fact, there's a command line port of the Unix Traceroute supplied with Windows 95 called TRACERT.EXE).

Traceroute uses ICMP (Internet Control Message Protocol) to attempt to trace the path an IP packet takes as it wends its way towards its intended host by initiating UDP probe packets with a TTL (Time To Live) which then listens for an ICMP "time exceeded" reply from a gateway.

Traceroutes start probing with a TTL of 1, increasing it by 1 until it receives an ICMP "port unreachable" error message which means Traceroute either got to the intended destination (the "host"), timed out, or hit a

"max" (in most implementations of Traceroute this defaults to 30 hops but can be changed by using the -m flag).

Three probes (which can also be modified by using the -q flag) are directed towards each TTL setting and the TTL address of the gateway and the round-trip time of each probe is returned.

If the Traceroute responses come from different gateways, the address of each responding system will be returned. However, if there's no response within a timeout interval (normally three seconds but this can be modified via the -w flag), an asterisk is displayed for that particular probe.

There are alternative ways to use ICMP packets to perform a very similar function, but ICMP normally only supports fewer than nine hops, which is a very small number for external Internet probes (although it may be perfectly suitable for internal network diagnosis). Traceroute's record route flag (or -r argument) to "ping" remote hosts is faster. It doesn't, however, return response statistics for the various legs of the journey and it is not totally reliable.

Results

Traceroute is a diagnostic tool that can safely be issued by your users when they are experiencing network problems (for example, when a particular Web page is unusually slow to download). Some early implementations of Traceroute did not work with certain combinations of Windows NT and specific (namely Cisco) router BIOS interfaces. These problems have now all been eliminated, except for a specific problem regarding TTL "shrouds" or packet-blocking daemons from Spammer sites that try to remain anonymous.

When a user is experiencing problems, it is best to ask them to use Traceroute from his or her terminal rather from your login or terminal (if you reside on a different part of the network). This is because you will need to eliminate router problems from within your own network before assuming the fault lies with someone else's routers outside the network.

Most of the common problems are delays caused by an overloaded server (too many users fighting for the same aggregated port address). The ICMP packets used by Traceroute are handled within the TCP/IP network code on the Web server. Because of this, the server's OS can still process ICMP efficiently - even if the server is overloaded with disk-read requests or is switching between too many server processes.

Many users assume that a request to download a Web page establishes some sort of two-way communication channel with its server. Traceroute's output reinforces that notion to many users by listing a series of routers as it "hops" between the sending workstation and its destination. However this is not exactly the case.

```
Tracing route to www.dancingfox.com [208.197.253.130] over a maximum of 30 hops:

  1   298 ms   327 ms   330 ms  ipt-fo2.proxy.aol.com [205.188.198.95]
  2   368 ms   438 ms   329 ms  tot-wj-dr4r.proxy.aol.com [205.188.198.124]
  3   362 ms   328 ms   382 ms  tpopd-rri1.red.aol.com [205.188.128.57]
  4   310 ms   329 ms   411 ms  tpopd-rre1-P5-0.red.aol.com [205.188.128.1]
  5   366 ms   331 ms   377 ms  pos12-0-0.gw1.iad2.alter.net [205.188.128.109]
  6   359 ms   321 ms   326 ms  133.ATM2-0.XR1.DCA1.ALTER.NET [146.188.161.194]
  7   347 ms   326 ms   382 ms  295.ATM3-0.TR1.DCA1.ALTER.NET [146.188.161.142]
  8   369 ms   476 ms   391 ms  101.ATM6-0.TR1.SCL1.ALTER.NET [146.188.136.222]
  9   356 ms   433 ms   382 ms  299.ATM7-0.XR1.SFO4.ALTER.NET [146.188.146.77]
 10   418 ms   434 ms   434 ms  191.ATM8-0-0.GW3.SFO1.ALTER.NET [146.188.145.229]
 11   409 ms   416 ms   399 ms  earthlink-oak-gw.customer.ALTER.NET [137.39.166.2]
 12   538 ms   549 ms   369 ms  208.197.248.18
 13   464 ms   420 ms   381 ms  www.dancingfox.com [208.197.253.130]

Trace complete.
```

Figure 2 - Output from TRACERT WWW.DANCINGFOX.COM.

Figure 3 - Cyberkit in action.

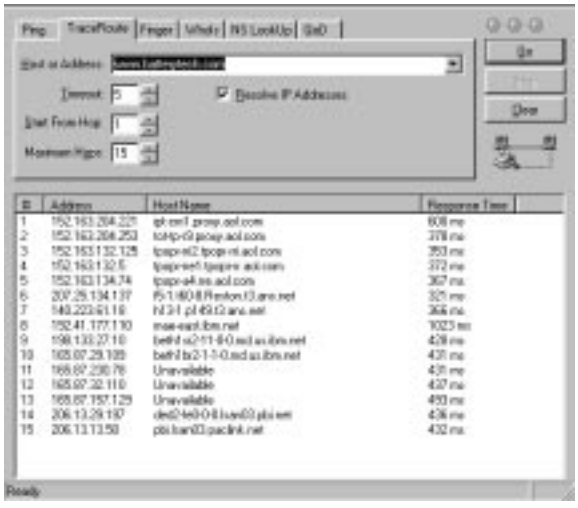

Figure 4 - Back with a different provider.

Traceroute

Traceroute probes the paths to a destination by sending ICMP bursts and incrementing the TTL on those packets. The first burst of three packets has a TTL of one hop, so they never get past the first router. Because these packets have used up their TTL count, the router will return a TE (or Time Exceeded) packet. Traceroute measures the delay between the original packet and the reply, which gives a very rough idea of how fast the connection is. Additional bursts of packets are sent with a TTL packet one hop greater than the previous TTL until the destination is reached (or until you reach the hop limit - which is 30 hops).

Because of this, it's not certain that the Traceroute replies will return via the same route through which the initial ICMP probes were initiated. Although routing issues were simpler in the early days of the Internet (and are still fairly straightforward over a LAN), it's now common that when using Traceroute to examine a server on the Internet - say 15 hops away - there can be as many as 15 different return routes used by each burst. If there are no replies received from one of the routers (which is indicated by a line of asterisks) it doesn't necessarily mean that the server is unreachable, only that the return path from this specific router is not usable. When you examine the Traceroute report, take care to check the return paths are valid.

Because a returned packet can often take a different path, this makes Traceroute unreliable for diagnosing certain problems involving in-bound packet losses and slow remote connections. Even if the server is unreachable, you cannot be entirely sure which router is at fault by reading the Traceroute report, because Traceroute can only return the outgoing path.

To understand how this works, here's an example.

We're at a company in London that gets its Internet connection from a fictional ISP called UKNet. We're accessing Web pages from a server in San Francisco belonging to a company that uses GenericNet.

The request travels from one of UKNet's local routers to a transatlantic gateway (from London to UKNet's main hub in Washington), where UKNet will then asymmetrically route the traffic onto GenericNet. GenericNet will then carry the request across to another address in Washington, and on to San Francisco. The Web server will then send us the page, which will travel on GenericNet's network back to Washington. From here it is switched back to UKNet via a "local hop" and then back to Washington, where it is passed across the Atlantic eventually ending up on our screens in London.

Asymmetric routing (sometimes called "hot potato" routing) is a common practice of backbone providers where any traffic on a provider's network that is destined to a point on another provider's network is handed over to the other provider at the earliest possible instant. This leads many users to assume (incorrectly) that problems lie with outside networks rather than their own. Because of this, many of the so-called intelligent diagnostic programs return non-existent problems which users call network support demanding a fix.

When To Use Traceroute

It is virtually impossible to accurately diagnose events on the Internet from a single location - especially if the complaints come from a computer connected via a dial-up modem connection, attempting to diagnose network congestion from an ISP's network or even a network link.

Congestion on a dial-up connection is often misinterpreted by diagnostic software as "backbone congestion" or as a problem with the ISP's router. Although it's possible this is the case, most ISPs will (rightly) refuse diagnostic information from dial-up users, simply because most ISPs will rely on their own internal network monitoring systems intrinsic to their routers and hardware. So if your users experience delays from remote connections dialling via ISPs with their local POPs (Points Of Presence) - themselves liable to congestion especially during peak periods when domestic users are clamouring for the same amount of aggregate bandwidth - then it's best to start looking at your own network before assuming a problem with the ISP's.

There are however a few useful results that can be gleaned from a dial-up PC with Ping and Traceroute. If a user reports a problem getting to a certain site, you can obtain the URL from them and Ping the domain name from the URL. This returns less data than Traceroute but it will tell you whether the Domain Name Server for the site is active, and whether the site is reach-

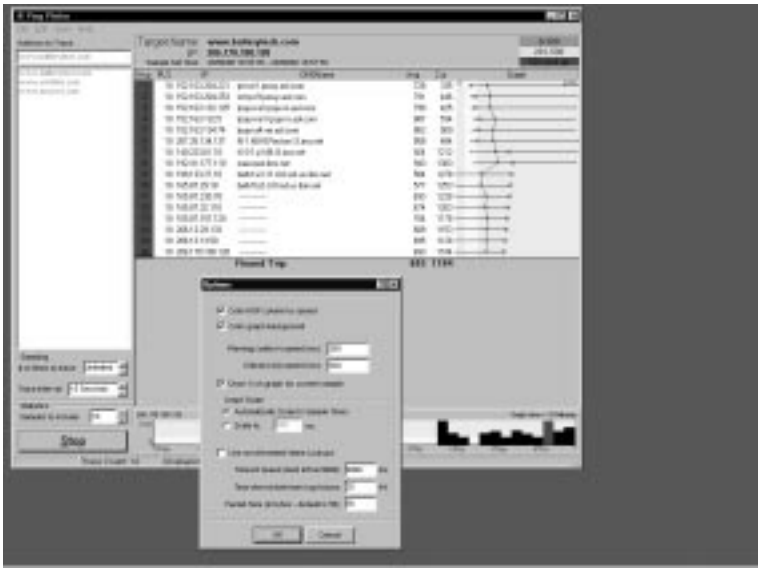

Figure 5 - PingPlotter.

able from the Internet. If the DNS is not up, there's nothing you can do except maybe contact the DNS's host or upstream provider. Such outages are normally due to system maintenance or hardware (router) failures and the upstream provider will probably know about it. If you're attempting to Ping a suspected spammer site, it may be that they have already been issued with a DoS (Denial Of Service) and their "nameservice" entry has been removed upstream.

However, if your search returns a domain name that resolves to an IP address, but you cannot successfully Ping that IP address, you can try using Traceroute to more accurately examine the outward journey. You might find out that at some point in the Traceroute, the turnaround time jumps significantly. You may also see that after a certain point you get a lot of asterisks in the listing. This indicates that packets have been significantly delayed somewhere along the round-trip.

Because you can't explore return paths via Traceroute, you can't accurately pinpoint where the problem lies. However, if you can decode the domain names at each hop, you have proof enough to inform your end-users that the problem really lies elsewhere on the remote network and not on your own network.

If Ping and Traceroute are both successful in negotiating a round trip - with figures to prove it - then your only recourse is to telnet to the problem site. From your user's PC, you can launch a Telnet session into a specific port. If this is a Web server, then you would typically Telnet into port 80 (although other ports are used for returning specific diagnostic information and messages) for remote access users.

The Hard Way

If in our example you want to access the Web server, type Telnet 80 ddd, where ddd is the domain name. Once successfully connected you can issue a GET / HTTP / 1.0 command (followed by two carriage returns). You should receive a stream of HTML code. If you see HTML code, then you've determined that the remote Web server is running and sending Web pages. If, however, you only get a stream of data beginning with "Date:" and "Content-type:" but there's no HTML code in sight, then it's likely that their server has a Maximum Transmission Unit (MTU) mismatch. There are utilities available that manually change the MTU setting on PCs, which may solve this problem (such as NetOptimizer and TweakDUN).

Automated Tools

Examining MTU settings and using Telnet to access a server's console or http port is network analysis and diagnosis at the very "nuts and bolts" level. Most system administrators will save themselves a lot of time and effort by setting up a local Web page which accesses a variety of diagnostic scripts that users can run from their own Web browsers. This takes time to set up and maintain and because of this, a small industry of diagnostic tools for Windows 95/98 and Windows NT has cropped up to make your life easier.

Built into every Windows 95 desktop is a Traceroute which operates in exactly the same way as the Unix command-line version. To run it, type TRACERT at the command line and you'll receive output rather like that shown in Figure 1.

To see how this all works, try tracerouting my Web site at www.dancingfox.com. The output is shown in Figure 2.

Cyberkit

Other Traceroutes exist, such as the one that forms part of Cyberkit, an extremely small postcardware (send the authors a postcard telling them how much you like it) package for Windows 95 / NT. It's a little friendlier than the Traceroute supplied with Windows 95/98 and it supplies a lot of alternative data about a server, its statistics as well as supply a range of other diagnostic tools such as Whois, MX (Mail eXchange) Records analysis and Finger.

Figure 3 shows Cyberkit in action. Here we have a situation where a Web server (www.batterytech.com) is not serving Web pages to customers. A Whois reveals that it is hosted at alter.net and analysing the path shows that, from AOL, it is just one jump from each other's networks (AOL.NET and alter.net) and that traffic between routers is reasonably healthy with no more than 485 milliseconds between each hop. It all progresses reasonably smoothly until it reaches hop 10, where a no-response is met.

Thus the network outage exists at some point within the alter.net network and not ours, and we can, if necessary, contact root@alter.net or webmaster@alter.net and inform them that there is some sort of outage starting at an internal ATM router, the address of which is 146.188.208.5.

In Figure 4, the server is backed up via a different provider and the routing topography has changed considerably. A later Traceroute shows that the Web server appears active now and the journey has complicated itself dramatically. Now jumping across to several sites, there is still a hefty wait of 1023 ms at ibm.net's east-coast hub.

PingPlotter

A very nice alternative to the text-based Traceroute packages and Cyberkit is the commercial PingPlotter (see Figure 5) - a fast, small and visual Traceroute utility. It uses multiple threads to trace all hops at once for substantial performance improvements over standard Traceroute packages - especially command line ones - and it draws a performance graph to pinpoint problems, helping you track a range of responses and trends. If a part of the network or a domain keeps giving you problems, you can include any number of samples in the graph with different sample times of your choosing (one second to many hours between samples) enabling you to log performance statistics. Both Cyberkit and PingPlotter can be downloaded on a trial basis from the Tucows Web site at www.tucows.com.

PCNA

The Author

Clive Grace (clive.grace@itp-journals.com) is a freelance IT writer.

New Reviews from Tech Support Alert

Anti-Trojan Software Reviews

A detailed review of six of the best anti trojan software programs. Two products were impressive with a clear gap between these and other contenders in their ability to detect and remove dangerous modern trojans.

Inkjet Printer Cartridge Suppliers

Everyone gets inundated by hundreds of ads for inkjet printer cartridges, all claiming to be the cheapest or best. But which vendor do you believe? Our editors decided to put them to the test by anonymously buying printer cartridges and testing them in our office inkjet printers. Many suppliers disappointed but we came up with several web sites that offer good quality cheap inkjet cartridges with impressive customer service.

Windows Backup Software

In this review we looked at 18 different backup software products for home or SOHO use. In the end we could only recommend six though only two were good enough to get our “Editor’s Choice” award

The 46 Best Freeware Programs

There are many free utilities that perform as well or better than expensive commercial products. Our Editor Ian Richards picks out his selection of the very best freeware programs and he comes up with some real gems.