

## Re: ICMP ping effecting network flow?

---

*Source:* <http://linux.derkeiler.com/Newsgroups/comp.os.linux.networking/2006-06/msg00424.html>

---

- *From:* Roy L Butler <[roy.butler@xxxxxxxxxxxxx](mailto:roy.butler@xxxxxxxxxxxxx)>
  - *Date:* Mon, 19 Jun 2006 10:30:55 -0700
- 

Moe Trin wrote:

On Fri, 16 Jun 2006, in the Usenet newsgroup comp.os.linux.networking, in article <e6ve31\$mu3\$1@xxxxxxxxxxxxxxxxxxxxxx>, Roy L Butler wrote:

I was told by someone at a systems conference about 10 years ago that if you ping a server over the Internet while downloading a file from it, your download will take place quicker.

I've come across documentation that some routers place ICMP traffic in their high-priority queue.

That still doesn't sound right. Lessee, 10 years ago... On 31 Dec 1995, the Internet consisted of 1.05 billion IP addresses assigned to 40865 networks world wide (for comparison, the figures earlier this week were 2.318 billion hosts and 73466 IPv4 networks) (Source: RIR Zone files dated 06/15/2006). If you look at any contemporary documentation (examples – the man page for the common LBL version of traceroute, or W. Richard Stevens "TCP/IP Illustrated Volume 1 – the protocols", Addison Wesley, ISBN 0-201-63346-9, copyright 1994) you'd see that many connections over the Internet require at least 6 hops (routers) to get anywhere – a more common value is over a dozen. In the traceroute man page, it took 18 hops to go from Lawrence Berkeley Labs (about 45 miles Southeast of Oakland) to MIT in Boston – and no less than 13 to get from LBL to a specific server in Berkeley (not more than 50 miles away). The Stevens book shows similar hop counts as well. It's a bit of a stretch to think that ten or fifteen intermediate routers would note a ping going by from "A" to "B", and therefore remember to pass TCP datagrams between those hosts with higher priorities. While not very likely at that era, it is also possible that pings could have been routed over different intermediate hops compared to TCP.

As far as prioritizing ICMP – can you give some cite? The IP header has

## Re: ICMP ping effecting network flow?

a TOS variable, and one of those is the "minimize delay" flag (see page 29 of RFC0791). These flags have in general been ignored, even back in the early 1990s (see the Stevens book). RFC1340 (superseded by RFC1700, which in turn was replaced by on-line databases) and RFC1349 (from July 1992) specified how these flags were to be set, and the 'Minimize Delay' flag was only set for 'Telnet/Rlogin', FTP control (but not data), TFTP, SMTP (command phase, not data) and DNS queries using UDP. These were generally interactive or setup conditions. There is also a 'Maximize throughput' flag which was set for bulk data transfers (FTP or SMTP data, DNS zone transfers).

Both these guys were heavily certified (which I know doesn't always mean qualified),

A lot depends on the certifying organization. Certifications are generally most useful in getting past the hiring filters set by HR and PHBs. The actual experience/knowledge would be determined in interviews with the technical people.

My guess is that they might have thought (wrongly?) that certain vendors' routers placed all traffic to/from a source address issuing echo\_requests in such a queue, elevating lower-priority protocols.

How much memory do you think is in the "average" backbone router? How much traffic do you think it might see in any given moment? (If you were traveling by air in the 60s-80s, you probably knew that if you traveled anywhere in the "South" of the United States (actually meaning Southeast), you couldn't go ANYWHERE without changing planes in Atlanta. By the mid-90s, the Internet was a good bit better than that, but there were still a lot of routes where the path from "A" to "B" went through major hubs like SFO, Denver, Chicago, and New York - even if you were going from Dallas to Atlanta.

I can't remember their names and I agree the time to ask was back then.

You might try 'comp.protocols.tcp-ip' or even 'alt.folklore.computers'. My news server also lists a 'alt.folklore.internet' which might be useful.

Old guy

Thanks for all of the good info. My specific source stating routers may be configured to treat ICMP echoes as high-priority packets is section 6.5 of Wide-Area Data Network Performance Engineering, by Cole & Ramaswamy (AT&T engineers), written in 2000. They do mention that the opposite (lower-priority) is common and I can understand how the whole thing seems even less plausible on a packet-switching

Re: ICMP ping effecting network flow?

architecture... I'll give up the position of devil's advocate. :)

Roy

.