

Re: icmp echo_user

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2007-10/msg00246.html>

- *From:* "Bruce M. Simpson" <bms@xxxxxxxxxxx>
 - *Date:* Sat, 27 Oct 2007 03:56:29 +0100
-

Matus Harvan wrote:

Hi,

I was wondering if I could get some feedback about the patch and whether others think it could be committed.

Thanks for your hard work on mtund. I'm not keen on this patch going into a mainline kernel, though.

It stomps on bandwidth limitation if that's in effect — which is a possible DoS vector — and also stops updating icmp protocol counters.

I believe we should track echo requests in netstat -p regardless of whether the kernel calls icmp_reflect() or not, as it can readily be inferred if a) your diversion to SOCK_RAW is in effect or b) the kernel processed the echo request.

I also believe that a user who installs and configures the tunneling daemon is in a position to know that the ICMP thresholds need to be changed.

Assuming the tunneling daemon doesn't process echoes unrelated to its tunneling (I haven't read the code), then the fact that rip_input() may exhaust its socket input buffer will provide a basic form of hysteresis, however I would suggest that if you intend to deploy this on the open Internet that the daemon either a) provides its own hysteresis too, b) tunes itself around the bandwidth limit in effect or c) tunes the bandwidth limit itself.

A better approach would be to conditionalise the 'goto raw' next to the 'goto reflect'.

regards,
BMS

freebsd-net@xxxxxxxxxxx mailing list
<http://lists.freebsd.org/mailman/listinfo/freebsd-net>
To unsubscribe, send any mail to "freebsd-net-unsubscribe@xxxxxxxxxxx"