

Re: UDP dont fragment bit

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2005-09/0183.html>

From: Sten Daniel Sørnsdal (*lists_at_wm-access.no*)

Date: 09/22/05

Date: Thu, 22 Sep 2005 13:34:10 +0200
To: Jeremie Le Hen <jeremie@le-hen.org>

Jeremie Le Hen wrote:

> *Hi,*
>
>
>
>> *Often there already is need for a tcp connection for authentication,*
>> *negotiation and so forth.*
>>
>> *RTT could, among other things, make a discovery process choose how fine*
>> *the increments/descrments should be.*
>>
>> *Estimated bandwidth could also help the actual data transport start out*
>> *with a more situation correct value.*
>>
>> *However with the abundance of routers modifying TCP MSS correctly*
>> *incorrectly and the odd chance that the data path of UDP packets is*
>> *different than TCP packets it wouldnt really give anything necessarily*
>> *reliable discovery process. And taking advantage of such values could*
>> *make the process more complex or less reliable.*
>
>
> *This is a rather specific case IMHO. BSD community didn't use to take*
> *care of such non standard behaviours, AFAIK. What you describe here*
> *is not really what's currently stated in RFCs.*

That it is not stated in an RFC is true. Please explain why it would have to be?

> *I would not be in favor of adding such an option in the FreeBSD kernel*
> *because, as Robert stated, this doesn't bring anything if not coupled*
> *with a non-trivial mechanim that could provide the user with ICMP MTU*
> *events. If one adds this option to the manual page, this will lead*
> *for sure to have mails emitted on this even list asking for how to*
> *retrieve those informations. Furthermore, I ought to add that the*
> *algorithm you described in pseudo-code lacks of robustness because*
> *of possible network congestion, which means this isn't something one*
> *would really see in wide use.*

freebsd-net: Re: UDP dont fragment bit

You assume such an application would need ICMP.

Again, let me state that such an application would not need ICMP's and it would work despite fragmentation issues.

Retrieve what information?

And an application based solely on my pseudo-code lacks robustness to network congestion is correct. So what?

- > *In other words, I think the feature you're calling for is really*
- > *specific to your problem, regarding your current network environment.*
- > *The misbehaviour of some particular network-fascist ISP should not*
- > *reach the FreeBSD source tree.*

You are mistaken in assuming my isp would block icmps or fragments. And you are mistaken that this was meant to solve a problem specific to one network. It is about giving unprivileged applications the opportunity to find the optimal packet size without relying on network policies or flaws.

--

Sten Daniel Sørdsal

freebsd-net@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-net>

To unsubscribe, send any mail to "freebsd-net-unsubscribe@freebsd.org"