

Re: 64 bit packet counters

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/net/2003-11/0053.html>

From: Jesper Skriver (jesper_at_FreeBSD.org)

Date: 11/09/03

Date: Sun, 9 Nov 2003 15:11:32 +0100

To: Harti Brandt <brandt@fokus.fraunhofer.de>

On Sat, Nov 08, 2003 at 10:13:48AM +0100, Harti Brandt wrote:

> *On Fri, 7 Nov 2003, Alex Hoff wrote:*

>

> *AH>Hi,*

> *AH>*

> *AH>We are attempting to implement the IF-MIB, which requires the*

> *AH>use of 64 bit packet counters and the differentiation between*

> *AH>multicast and broadcast pkts. Since changing the if_data (by*

> *AH>adding new counters and changing the existing to u_int64) is a bad*

> *AH>idea, does anyone have any good ideas on how to do this? I was*

> *AH>thinking of tacking on a new struct (lets call it ifx_data) on at*

> *AH>the end of the current if_net struct with the appropriate counters*

> *AH>(i/opacket, i/obyte, i/obcast, i/omcast). Apart from having to do*

> *AH>a little double counting is there any obvious pitfals with this*

> *AH>approach? Does anyone have an better ideas? Is there currently any*

> *AH>plans to update the network stack to handle this properly?*

>

> *You may lookup the discussions in the mailing lists. As far as*

> *I remember the problem with 64 bit counting was that this needs*

> *locks because not on all architectures you have atomic 64bit add*

> *operations. A simple method that does not involve kernel changes (and*

> *that I plan to implement in my snmp daemon) is to periodically monitor*

> *the counters (depending on the interface speed) and detect wraps in*

> *the daemon.*

What is done in other places is to have normal 32 bit counters that are incremented per packet, and have a background task/thread to update the 64 bit counters from the 32 bit counters.

That way, we avoid the locking issue per packet.

/Jesper

--

Jesper Skriver, [jesper\(at\)skriver\(dot\)dk](mailto:jesper(at)skriver(dot)dk) - CCIE #5456
One Unix to rule them all, One Resolver to find them,
One IP to bring them all and in the zone to bind them.

freebsd-net: Re: 64 bit packet counters

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"