Re: [RFC][UPDATED PATCH 2.6.16] [Patch 9/9] Generic netlink interface for delay accounting

Source: http://linux.derkeiler.com/Mailing−Lists/Kernel/2006−03/msg08517.html

- *From*: Balbir Singh <balbir@xxxxxxxxxx>
- *Date*: Fri, 24 Mar 2006 20:24:59 +0530

On Fri, Mar 24, 2006 at 09:04:21AM −0500, jamal wrote:

> On Thu, 2006−23−03 at 21:11 +0530, Balbir Singh wrote:
>
> > On Thu, Mar 23, 2006 at 09:04:46AM −0500, jamal wrote:
> >
> > > Should there be at least either a pid or tgid? If yes, you need
> > > to
> > > validate here...
> >
> > Yes, you are correct. One of my test cases caught it too.. But I did
> > not want to untidy the code with if−else's which will keep growing if
> > the attributes change in the future. I just followed the controller
> > example. I will change it and validate it. Currently if the attribute
> > is not valid, a stat of all zero's is returned back.
>
> There are many ways to skin this cat.
> As an example: you could make pid and tgid global to the function and
> set them to 0. At the end of the if statements, you could check if at
> least one of them is set − if not you know none was passed and bail out.

The latest patch does fix it this issue. In the Changelog
6. taskstats_send_stats() now validates the command attributes and ensures
that it either gets a PID or a TGID. If it gets both simultaneously
the PID stats are sent.

Is this change ok with you?

> As a general comment double check your logic for errors; if you already
> have stashed something in the skb, you need to remove it etc.

> Wouldn't genlmsg_cancel() take care of cleaning all attributes?

> it would for attribute setting − but not for others. As a general
> comment this is one of those areas where cutnpasting aka TheLinuxWay(tm)
> could result in errors.

:−) I understand.
What I have done is moved all the NLA_PUT_U32 to after verifying the return values of functions fill_*(). That way we do not stash anything into the skb if there are pending errors.

> A single message with PID+TGID sounds reasonable. Why two messages with
> two stats? all you will need to do is get rid of the prepare_reply()
> above and NLA_PUT_U32() below (just like you do in a response to a GET.

> The reason for two stats is that for TGID, we return accumulated values
> (of all threads in the group) and for PID we return the value just
> for that pid. The return value is

> Ok, I understand the dilemma now − but still not thrilled with having
> two messages.
> Perhaps you could have nesting of TLVs? This is widely used in the net
> code for example
> i.e:
>
> TLV = TASKSTATS_TYPE_TGID/PID
> TLV = TASKSTATS_TYPE_STATS
>
> Look at using nla_nest_start/end/cancel

Hmm... Would it be ok to send one message with the following format

1. TLV=TASKSTATS_TYPE_PID
2. TLV=TASKSTATS_TYPE_STATS
3. TLV=TASKSTATS_TYPE_TGID
4. TLV=TASKSTATS_TYPE_STATS

It would still be one message, except that 3 and 4 would be optional.
What do you think?

cheers,
jamal

Thanks for your comments,
Balbir
−
To unsubscribe from this list: send the line "unsubscribe linux−kernel" in
the body of a message to majordomo@xxxxxxxxxxxxxxx
More majordomo info at http://vger.kernel.org/majordomo−info.html
Please read the FAQ at http://www.tux.org/lkml/