
Subject: Re: [PATCH 1/3] lockd: use rpc client's cl_nodename for id encoding
Posted by [Chuck Lever](#) on Fri, 14 Sep 2012 19:08:42 GMT
[View Forum Message](#) <> [Reply to Message](#)

Hi-

On Sep 14, 2012, at 10:25 AM, Stanislav Kinsbursky wrote:

```
> Taking hostname from uts namespace if not safe, because this could be
> performind during umount operation on child reaper death. And in this case
> current->nsproxy is NULL already.
>
> Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>
> Cc: <stable@vger.kernel.org>
> ---
> fs/lockd/mon.c | 14 ++++++-----
> 1 files changed, 8 insertions(+), 6 deletions(-)
>
> diff --git a/fs/lockd/mon.c b/fs/lockd/mon.c
> index 7ef14b3..c6186fb 100644
> --- a/fs/lockd/mon.c
> +++ b/fs/lockd/mon.c
> @@ -426,11 +426,12 @@ static void encode_mon_name(struct xdr_stream *xdr, const struct
nsm_args *argp)
> * (via the NLMPROC_SM_NOTIFY call) that the state of host "mon_name"
> * has changed.
> */
> -static void encode_my_id(struct xdr_stream *xdr, const struct nsm_args *argp)
> +static void encode_my_id(struct xdr_stream *xdr, const struct nsm_args *argp,
> + char *nodename)
> {
>     __be32 *p;
>
> - encode_nsm_string(xdr, utsname()->nodename);
> + encode_nsm_string(xdr, nodename);
> p = xdr_reserve_space(xdr, 4 + 4 + 4);
> *p++ = cpu_to_be32(argp->prog);
> *p++ = cpu_to_be32(argp->vers);
> @@ -441,10 +442,11 @@ static void encode_my_id(struct xdr_stream *xdr, const struct
nsm_args *argp)
> * The "mon_id" argument specifies the non-private arguments
> * of an NSMPROC_MON or NSMPROC_UNMON call.
> */
> -static void encode_mon_id(struct xdr_stream *xdr, const struct nsm_args *argp)
> +static void encode_mon_id(struct xdr_stream *xdr, const struct nsm_args *argp,
> + char *nodename)
> {
>     encode_mon_name(xdr, argp);
```

```

> - encode_my_id(xdr, argp);
> + encode_my_id(xdr, argp, nodename);
> }
>
> /*
> @@ -463,14 +465,14 @@ static void encode_priv(struct xdr_stream *xdr, const struct
nsm_args *argp)
> static void nsm_xdr_enc_mon(struct rpc_rqst *req, struct xdr_stream *xdr,
>     const struct nsm_args *argp)
> {
> - encode_mon_id(xdr, argp);
> + encode_mon_id(xdr, argp, req->rq_task->tk_client->cl_nodename);

```

IMO you should get the cl_nodename in nsm_mon_unmon() from clnt->cl_nodename, and pass it in as part of *argp . The choice of which nodename to use is clearly a decision for an "upper layer" not a choice for the XDR functions.

Long ago I had patches which fixed this layering violation for a very similar purpose as yours, but they were never applied.

```

> encode_priv(xdr, argp);
> }
>
> static void nsm_xdr_enc_unmon(struct rpc_rqst *req, struct xdr_stream *xdr,
>     const struct nsm_args *argp)
> {
> - encode_mon_id(xdr, argp);
> + encode_mon_id(xdr, argp, req->rq_task->tk_client->cl_nodename);

```

Ditto.

```

> }
>
> static int nsm_xdr_dec_stat_res(struct rpc_rqst *rqstp,
>
> --
> To unsubscribe from this list: send the line "unsubscribe linux-nfs" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html

```

--
Chuck Lever
chuck[dot]lever[at]oracle[dot]com