On Thu, Jan 19, 2012 at 09:04:40PM +0400, Stanislav Kinsbursky wrote:

> >On Thu, Jan 19, 2012 at 06:49:23PM +0400, Stanislav Kinsbursky wrote:
> >>@@ -1000,6 +996,7 @@ static int svcauth_gss_handle_init(struct svc_rqst *rqstp,
> >>   struct xdr_netobj tmpobj;
> >>   struct rsi *rsip, rsikey;
> >>   int ret;
> >>+ struct sunrpc_net *sn = net_generic(rqstp->rq_xprt->xpt_net, sunrpc_net_id);
> >
> >OK, so you're getting the network namespace out of the rqstp, and, then
> >passing it down, makes sense. And:
> >
> >>@@ -1079,6 +1076,7 @@ svcauth_gss_accept(struct svc_rqst *rqstp, __be32 *authp)
> >>   __be32 *rpcstart;
> >>   __be32 *reject_stat = resv->iov_base + resv->iov_len;
> >>   int ret;
> >>+ struct sunrpc_net *sn = net_generic(rqstp->rq_xprt->xpt_net, sunrpc_net_id);
> >
> >... same for the gss cache. Looks good.
> >
> >How do you plan to test this?
> >
>
> Do you mean something special or in general?
> Currently I validate all my chages in container by using simple test environment.
> I would be appreciate for any hints to tests, than can help.

The server needs to be tested after these changes, and we need to make
sure the caches affected still work.

(I suspect the gid cache code will oops if it's used after these
patches, since the table is left NULl?)

And then we should also test in a container environment, with different
instances of mountd and rpc.svcidmapd running in each container, to
verify that the right thing happens.

--b.