
Subject: [PATCH] SUNRPC: skip dead but not buried clients on PipeFS events
Posted by [Stanislav Kinsbursky](#) on Thu, 19 Apr 2012 10:57:18 GMT

[View Forum Message](#) <> [Reply to Message](#)

These clients can't be safely dereferenced if their counter in 0.

Signee-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

Signed-off-by: Stanislav Kinsbursky <skinsbursky@parallels.com>

net/sunrpc/clnt.c | 3 ++-
1 files changed, 2 insertions(+), 1 deletions(-)

diff --git a/net/sunrpc/clnt.c b/net/sunrpc/clnt.c
index 6797246..591994d 100644
--- a/net/sunrpc/clnt.c
+++ b/net/sunrpc/clnt.c
@@ -218,7 +218,8 @@ static struct rpc_clnt *rpc_get_client_for_event(struct net *net, int event)
if (((event == RPC_PIPEFS_MOUNT) && clnt->cl_dentry) ||
((event == RPC_PIPEFS_UMOUNT) && !clnt->cl_dentry))
continue;
- atomic_inc(&clnt->cl_count);
+ if (atomic_inc_return(&clnt->cl_count) == 1)
+ continue;
spin_unlock(&sn->rpc_client_lock);
return clnt;
}

Subject: Re: [PATCH] SUNRPC: skip dead but not buried clients on PipeFS events
Posted by [Stanislav Kinsbursky](#) on Thu, 19 Apr 2012 11:36:57 GMT

[View Forum Message](#) <> [Reply to Message](#)

Sorry, but ignore this patch too.

It can't be that simple because of these cl_count tricks in rpc_release_client...

> These clients can't be safely dereferenced if their counter in 0.
>
> Signee-off-by: Stanislav Kinsbursky<skinsbursky@parallels.com>
>
> Signed-off-by: Stanislav Kinsbursky<skinsbursky@parallels.com>
>
> ---
> net/sunrpc/clnt.c | 3 ++-
> 1 files changed, 2 insertions(+), 1 deletions(-)
>

```
> diff --git a/net/sunrpc/clnt.c b/net/sunrpc/clnt.c
> index 6797246..591994d 100644
> --- a/net/sunrpc/clnt.c
> +++ b/net/sunrpc/clnt.c
> @@ -218,7 +218,8 @@ static struct rpc_clnt *rpc_get_client_for_event(struct net *net, int
event)
>     if (((event == RPC_PIPEFS_MOUNT)&& clnt->cl_dentry) ||
>         ((event == RPC_PIPEFS_UMOUNT)&& !clnt->cl_dentry))
>         continue;
> - atomic_inc(&clnt->cl_count);
> + if (atomic_inc_return(&clnt->cl_count) == 1)
> + continue;
> spin_unlock(&sn->rpc_client_lock);
> return clnt;
> }
>
```

--

Best regards,
Stanislav Kinsbursky

Subject: Re: [PATCH] SUNRPC: skip dead but not buried clients on PipeFS events
Posted by [bfields](#) on Thu, 19 Apr 2012 21:40:43 GMT

[View Forum Message](#) <> [Reply to Message](#)

On Thu, Apr 19, 2012 at 03:36:57PM +0400, Stanislav Kinsbursky wrote:

```
> Sorry, but ignore this patch too.
> It can't be that simple because of these cl_count tricks in rpc_release_client...
```

OK. Very minor whine:

```
>
```

```
> > These clients can't be safely dereferenced if their counter is 0.
> >
> > Signed-off-by: Stanislav Kinsbursky<skinsbursky@parallels.com>
> >
> > Signed-off-by: Stanislav Kinsbursky<skinsbursky@parallels.com>
```

I don't mind fixing up trivial slips every now and then, but that double signed-off-by seems to happen on a lot of your posts; could you figure out what's up with your scripts?

--b.

```
> >
```

```

> >---
> > net/sunrpc/clnt.c | 3 +-
> > 1 files changed, 2 insertions(+), 1 deletions(-)
> >
> >diff --git a/net/sunrpc/clnt.c b/net/sunrpc/clnt.c
> >index 6797246..591994d 100644
> >--- a/net/sunrpc/clnt.c
> >+++ b/net/sunrpc/clnt.c
> >@@ -218,7 +218,8 @@ static struct rpc_clnt *rpc_get_client_for_event(struct net *net, int
event)
> > if (((event == RPC_PIPEFS_MOUNT)&& clnt->cl_dentry) ||
> >      ((event == RPC_PIPEFS_UMOUNT)&& !clnt->cl_dentry))
> >     continue;
> >- atomic_inc(&clnt->cl_count);
> >+ if (atomic_inc_return(&clnt->cl_count) == 1)
> >+     continue;
> > spin_unlock(&sn->rpc_client_lock);
> > return clnt;
> > }
> >
> >
> >
> --
> Best regards,
> Stanislav Kinsbursky
> --
> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html
> Please read the FAQ at http://www.tux.org/lkml/

```

Subject: Re: [PATCH] SUNRPC: skip dead but not buried clients on PipeFS events
Posted by [Stanislav Kinsbursky](#) on Fri, 20 Apr 2012 06:05:05 GMT
[View Forum Message](#) <> [Reply to Message](#)

```

> On Thu, Apr 19, 2012 at 03:36:57PM +0400, Stanislav Kinsbursky wrote:
>> Sorry, but ignore this patch too.
>> It can't be that simple because of these cl_count tricks in rpc_release_client...
> OK. Very minor whine:
>

```

```

>>> These clients can't be safely dereferenced if their counter in 0.
>>>
>>> Signee-off-by: Stanislav Kinsbursky<skinsbursky@parallels.com>
>>>
>>> Signed-off-by: Stanislav Kinsbursky<skinsbursky@parallels.com>

```

> I don't mind fixing up trivial slips every now and then, but that double
> signed-off-by seems to happen on a lot of your posts; could you figure
> out what's up with your scripts?

Yes, sure. Sorry.

I've noticed this problem just yesterday.
