

---

Subject: Re: [PATCH RFC 00/13] Lockd: grace period containerization  
Posted by [bfields](#) on Wed, 16 May 2012 20:54:58 GMT

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

---

On Wed, May 16, 2012 at 04:06:30PM -0400, J. Bruce Fields wrote:

> On Sat, May 05, 2012 at 09:21:30PM +0400, Stanislav Kinsbursky wrote:  
> > This patch set is marked with RFC, because I'm still not quite sure, that this  
> > implementation will satisfy other interested people.  
> > So, would be appreciated for any comments.  
> >  
> > This patch set makes grace period and hosts reclaiming network namespace  
> > aware.  
> >  
> > Main ideas:  
> > 1) moving of  
> >  
> > unsigned long next\_gc;  
> > unsigned long nrhosts;  
> >  
> > struct delayed\_work grace\_period\_end;  
> > struct lock\_manager lockd\_manager;  
> > struct list\_head grace\_list;  
> >  
> > to per-net Lockd data.  
> >  
> > 2) moving of  
> >  
> > struct lock\_manager nfsd4\_manager;  
> >  
> > to per-net NFSd data.  
> >  
> > 3) shutdown + gc of NLM hosts done now network namespace aware.  
>  
> That all sounds reasonable to me.  
>  
> > 4) restart\_grace() now works only for init\_net.  
>  
> Eventually we might just remove that. I doubt it's used anywhere.

And on a quick skim I don't see anything wrong with the patches themselves; let me know when you consider them ready.

The per-net grace period management probably isn't what we'll want eventually, but as I said on the other thread I think it's a reasonable starting point.

--b.

---