
Subject: Re: Network namespaces a path to mergable code.
Posted by [Cedric Le Goater](#) on Wed, 28 Jun 2006 09:54:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

Eric W. Biederman wrote:

> Despite what it might look like unix domain sockets do not live in the
> filesystem. They store a cookie in the filesystem that roughly
> corresponds to the port number of an AF_INET socket. When you open a
> socket the lookup is done by the cookie retrieved from the filesystem.

unix domain socket lookup uses a path_lookup for sockets in the filesystem
namespace and a find_by_name for socket in the abstract namespace.

> So except for their cookies unix domain sockets are always in the
> network stack.

what is that cookie ? the file dentry and mnt ref ?

so, ok, the resulting struct sock is part of the network namespace but
there is a bridge with the filesystem namespace which does not prevent
other namespaces to do a lookup. the lookup routine needs to be changed,
this is any way necessary for the abstract namespace.

I think we're reaching the limits of namespaces. It would be much easier
with a container id in each kernel object we want to isolate.

C.
