

---

Subject: Re: [PATCH 4/9] network namespaces: socket hashes  
Posted by [Andrey Savochkin](#) on Wed, 20 Sep 2006 16:32:22 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Hi,

On Mon, Sep 18, 2006 at 05:12:49PM +0200, Daniel Lezcano wrote:

> Andrey Savochkin wrote:

> > Socket hash lookups are made within namespace.

> > Hash tables are common for all namespaces, with

> > additional permutation of indexes.

>

> Hi Andrey,

>

> why is the hash table common and not instanciated multiple times for

> each namespace like the routes ?

The main reason is that socket hash tables should be large enough to work efficiently, but it isn't good to waste a lot of memory for each namespace. Namespaces should be cheap enough, to allow to have hundreds of them. This reason of memory efficiency, of course, has a priority unless/until socket hash tables start to resize automatically.

Another point is that routing lookup is much more complicated than the socket's one to add another search key.

Routing also have additional routines for deleting entries matching some patterns, and so on.

In short, routing is much more complicated, and it already quite efficient for various sizes of routing tables.

Andrey

---