

---

Subject: Re: [PATCH 5/12] L2 network namespace: IPv4 routing

Posted by [Daniel Lezcano](#) on Wed, 06 Dec 2006 22:40:29 GMT

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

---

Dmitry Mishin wrote:

> Make FIBs per-namespace and adds additional key (net namespace) to lookups in  
> routing cache.

>

> Signed-off-by: Dmitry Mishin <dim@openvz.org>

>

> ---

> include/linux/net\_namespace.h | 12 ++++  
> include/net/flow.h | 3 +  
> include/net/ip\_fib.h | 48 ++++++++-----  
> net/core/fib\_rules.c | 43 ++++++++-----  
> net/core/net\_namespace.c | 9 +++  
> net/ipv4/fib\_frontend.c | 121 ++++++++-----  
> net/ipv4/fib\_hash.c | 13 +++-  
> net/ipv4/fib\_rules.c | 86 ++++++++-----  
> net/ipv4/fib\_semantics.c | 100 ++++++++-----  
> net/ipv4/fib\_trie.c | 18 ++++--  
> net/ipv4/route.c | 29 ++++++++--  
> 11 files changed, 391 insertions(+), 91 deletions(-)

>

> --- linux-2.6.19-rc6-mm2.orig/include/linux/net\_namespace.h

> +++ linux-2.6.19-rc6-mm2/include/linux/net\_namespace.h

> @@ -11,6 +11,18 @@ struct net\_namespace {  
> struct net\_device \*dev\_base\_p, \*\*dev\_tail\_p;  
> struct net\_device \*loopback\_dev\_p;  
> struct pcpu\_lstats \*pcpu\_lstats\_p;  
> +#ifndef CONFIG\_IP\_MULTIPLE\_TABLES  
> + struct fib\_table \*fib4\_local\_table, \*fib4\_main\_table;  
> +#else  
> + struct list\_head fib\_rules\_ops\_list;  
> + struct fib\_rules\_ops \*fib4\_rules\_ops;  
> +#endif  
> + struct hlist\_head \*fib4\_tables;  
> + struct hlist\_head \*fib4\_hash, \*fib4\_laddrhash;  
> + unsigned fib4\_hash\_size, fib4\_info\_cnt;  
> +#ifdef CONFIG\_IP\_FIB\_TRIE  
> + int fib4\_trie\_last\_dflt;  
> +#endif  
> unsigned int hash;  
> };  
>  
> --- linux-2.6.19-rc6-mm2.orig/include/net/flow.h  
> +++ linux-2.6.19-rc6-mm2/include/net/flow.h  
> @@ -82,6 +82,9 @@ struct flowi {

```

> #define fl_mh_type uli_u.mht.type
> #endif
> __u32 secid; /* used by xfrm; see secid.txt */
> #ifdef CONFIG_NET_NS
> + struct net_namespace *net_ns;
> #endif
> } __attribute__((__aligned__(BITS_PER_LONG/8)));
>
> #define FLOW_DIR_IN 0
> --- linux-2.6.19-rc6-mm2.orig/include/net/ip_fib.h
> +++ linux-2.6.19-rc6-mm2/include/net/ip_fib.h
> @@ -18,6 +18,7 @@
>
> #include <net/flow.h>
> #include <linux/seq_file.h>
> #include <linux/net_namespace.h>
> #include <net/fib_rules.h>
>
> struct fib_config {
> @@ -171,14 +172,21 @@ struct fib_table {
>
> #ifndef CONFIG_IP_MULTIPLE_TABLES
>
> -extern struct fib_table *ip_fib_local_table;
> -extern struct fib_table *ip_fib_main_table;
> #ifdef CONFIG_NET_NS
> +extern struct fib_table *ip_fib_local_table_static;
> +extern struct fib_table *ip_fib_main_table_static;
> #define ip_fib_local_table_ns() ip_fib_local_table_static
> #define ip_fib_main_table_ns() ip_fib_main_table_static
> #else
> #define ip_fib_local_table_ns() (current_net_ns->fib4_local_table)
> #define ip_fib_main_table_ns() (current_net_ns->fib4_main_table)
> #endif

```

Why don't you use :

```

#ifndef CONFIG_NET_NS
extern struct fib_table *ip_fib_local_table
extern struct fib_table *ip_fib_main_table
#else
#define ip_fib_local_table (current_net_ns->fib4_local_table)
#define ip_fib_main_table (current_net_ns->fib4_main_table)
#endif

```

You don't need to replace all `ip_fib_local_table` by `ip_fib_local_table_ns`. It will reduce a lot the code impact, no ?

---

Containers mailing list  
Containers@lists.osdl.org  
<https://lists.osdl.org/mailman/listinfo/containers>

---