

---

Subject: [PATCH][net-2.6.25][NETNS] make netns cleanup to run in a separate workqueue

Posted by [Daniel Lezcano](#) on Mon, 19 Nov 2007 10:44:25 GMT

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

---

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

---

Subject: Re: [PATCH][net-2.6.25][NETNS] make netns cleanup to run in a separate workqueue

Posted by [dev](#) on Mon, 19 Nov 2007 11:01:27 GMT

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

---

imho panic() is too much.

create\_singlethread\_workqueue() can fail e.g. due to out of memory...

Thanks,  
Kirill

Daniel Lezcano wrote:

> Subject: make netns cleanup to run in a separate queue

> From: Benjamin Thery <benjamin.thery@bull.net>

>

> This patch adds a separate workqueue for cleaning up a network  
> namespace. If we use the keventd workqueue to execute cleanup\_net(),  
> there is a problem to unregister devices in IPv6. Indeed the code  
> that cleans up also schedule work in keventd: as long as cleanup\_net()  
> hasn't return, dst\_gc\_task() cannot run and as long as dst\_gc\_task() has  
> not run, there are still some references pending on the net devices and  
> cleanup\_net() can not unregister and exit the keventd workqueue.

>

> Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>

> Signed-off-by: Benjamin Thery <benjamin.thery@bull.net>

> Acked-by: Denis V. Lunev <den@openvz.org>

> ---

> net/core/net\_namespace.c | 9 ++++++++

> 1 file changed, 8 insertions(+), 1 deletion(-)

>

> Index: net-2.6.25/net/core/net\_namespace.c

> =====

> --- net-2.6.25.orig/net/core/net\_namespace.c

> +++ net-2.6.25/net/core/net\_namespace.c

> @@ -58,6 +58,7 @@ out\_undo:

>

```

> #ifdef CONFIG_NET_NS
> static struct kmem_cache *net_cachep;
> +static struct workqueue_struct *netns_wq;
>
> static struct net *net_alloc(void)
> {
> @@ -149,7 +150,7 @@ void __put_net(struct net *net)
> {
> /* Cleanup the network namespace in process context */
> INIT_WORK(&net->work, cleanup_net);
> - schedule_work(&net->work);
> + queue_work(netns_wq, &net->work);
> }
> EXPORT_SYMBOL_GPL(__put_net);
>
> @@ -171,7 +172,13 @@ static int __init net_ns_init(void)
> net_cachep = kmem_cache_create("net_namespace", sizeof(struct net),
> SMP_CACHE_BYTES,
> SLAB_PANIC, NULL);
> +
> + /* Create workqueue for cleanup */
> + netns_wq = create_singlethread_workqueue("netns");
> + if (!netns_wq)
> + panic("Could not create netns workq");

> #endif
> +
> mutex_lock(&net_mutex);
> err = setup_net(&init_net);
>
>
>
> -----
>
> _____
> Containers mailing list
> Containers@lists.linux-foundation.org
> https://lists.linux-foundation.org/mailman/listinfo/containers

```

---

Containers mailing list  
Containers@lists.linux-foundation.org  
<https://lists.linux-foundation.org/mailman/listinfo/containers>

---

Subject: Re: [PATCH][net-2.6.25][NETNS] make netns cleanup to run in a separate workqueue

Posted by [dev](#) on Mon, 19 Nov 2007 11:07:50 GMT

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

---

Ah, sorry. Didn't notice it's called only on boot.

Acked-By: Kirill Korotaev <dev@sw.ru>

Kirill Korotaev wrote:

> imho panic() is too much.

> create\_singlethread\_workqueue() can fail e.g. due to out of memory...

>

> Thanks,

> Kirill

>

>

> Daniel Lezcano wrote:

>

>>Subject: make netns cleanup to run in a separate queue

>>From: Benjamin Thery <benjamin.thery@bull.net>

>>

>>This patch adds a separate workqueue for cleaning up a network  
>>namespace. If we use the keventd workqueue to execute cleanup\_net(),  
>>there is a problem to unregister devices in IPv6. Indeed the code  
>>that cleans up also schedule work in keventd: as long as cleanup\_net()  
>>hasn't return, dst\_gc\_task() cannot run and as long as dst\_gc\_task() has  
>>not run, there are still some references pending on the net devices and  
>>cleanup\_net() can not unregister and exit the keventd workqueue.

>>

>>Signed-off-by: Daniel Lezcano <dlezcano@fr.ibm.com>

>>Signed-off-by: Benjamin Thery <benjamin.thery@bull.net>

>>Acked-by: Denis V. Lunev <den@openvz.org>

>>---

>> net/core/net\_namespace.c | 9 +++++++--

>> 1 file changed, 8 insertions(+), 1 deletion(-)

>>

>>Index: net-2.6.25/net/core/net\_namespace.c

>>=====

>>--- net-2.6.25.orig/net/core/net\_namespace.c

>>+++ net-2.6.25/net/core/net\_namespace.c

>>@@ -58,6 +58,7 @@ out\_undo:

>>

>> #ifdef CONFIG\_NET\_NS

>> static struct kmem\_cache \*net\_cachep;

>>+static struct workqueue\_struct \*netns\_wq;

>>

>> static struct net \*net\_alloc(void)

>> {

>>@@ -149,7 +150,7 @@ void \_\_put\_net(struct net \*net)

>> {

```

>> /* Cleanup the network namespace in process context */
>> INIT_WORK(&net->work, cleanup_net);
>>- schedule_work(&net->work);
>>+ queue_work(netns_wq, &net->work);
>> }
>> EXPORT_SYMBOL_GPL(__put_net);
>>
>>@@ -171,7 +172,13 @@ static int __init net_ns_init(void)
>> net_cachep = kmem_cache_create("net_namespace", sizeof(struct net),
>>   SMP_CACHE_BYTES,
>>   SLAB_PANIC, NULL);
>>+
>>+ /* Create workqueue for cleanup */
>>+ netns_wq = create_singlethread_workqueue("netns");
>>+ if (!netns_wq)
>>+   panic("Could not create netns workq");
>
>
>> #endif
>>+
>> mutex_lock(&net_mutex);
>> err = setup_net(&init_net);
>>
>>
>>-----
>>
>>_____
>>Containers mailing list
>>Containers@lists.linux-foundation.org
>>https://lists.linux-foundation.org/mailman/listinfo/containers
>
>
>_____
> Containers mailing list
> Containers@lists.linux-foundation.org
> https://lists.linux-foundation.org/mailman/listinfo/containers
>

```

---

```

Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers

```

---



---

Subject: Re: [PATCH][net-2.6.25][NETNS] make netns cleanup to run in a separate workqueue

Posted by [davem](#) on Tue, 20 Nov 2007 07:18:34 GMT

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

---

From: Kirill Korotaev <dev@sw.ru>

Date: Mon, 19 Nov 2007 14:08:49 +0300

> Acked-By: Kirill Korotaev <dev@sw.ru>

I've applied this patch, thanks everyone.

---

Containers mailing list

[Containers@lists.linux-foundation.org](mailto:Containers@lists.linux-foundation.org)

<https://lists.linux-foundation.org/mailman/listinfo/containers>

---