
Subject: [PATCH] remove __net_initdata attribute
Posted by [den](#) on Fri, 26 Oct 2007 09:54:26 GMT

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

The discussion with Eric W. Biederman reveals that the concept for __net_initdata is bogus. The ->exit method of pernet_operations should be called during module exit.

So, discard this attribute.

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

```
diff --git a/drivers/net/loopback.c b/drivers/net/loopback.c
index 662b8d1..45f30a2 100644
--- a/drivers/net/loopback.c
+++ b/drivers/net/loopback.c
@@ -284,7 +284,7 @@ static __net_exit void loopback_net_exit(struct net *net)
    unregister_netdev(dev);
}

-static struct pernet_operations __net_initdata loopback_net_ops = {
+static struct pernet_operations loopback_net_ops = {
    .init = loopback_net_init,
    .exit = loopback_net_exit,
};

diff --git a/fs/proc/proc_net.c b/fs/proc/proc_net.c
index 2e91fb7..cbc5a95 100644
--- a/fs/proc/proc_net.c
+++ b/fs/proc/proc_net.c
@@ -185,7 +185,7 @@ static __net_exit void proc_net_ns_exit(struct net *net)
    kfree(net->proc_net_root);
}

-struct pernet_operations __net_initdata proc_net_ns_ops = {
+struct pernet_operations proc_net_ns_ops = {
    .init = proc_net_ns_init,
    .exit = proc_net_ns_exit,
};

diff --git a/include/net/net_namespace.h b/include/net/net_namespace.h
index 93aa87d..5279466 100644
--- a/include/net/net_namespace.h
+++ b/include/net/net_namespace.h
@@ -102,11 +102,9 @@ static inline void release_net(struct net *net)
#endif CONFIG_NET_NS
#define __net_init
#define __net_exit
-#define __net_initdata
#else
```

```

#define __net_init __init
#define __net_exit __exit_refok
#define __net_initdata __initdata
#endif

struct pernet_operations {
diff --git a/net/core/dev.c b/net/core/dev.c
index f1647d7..1afcb56 100644
--- a/net/core/dev.c
+++ b/net/core/dev.c
@@ -2665,7 +2665,7 @@ static void __net_exit dev_proc_net_exit(struct net *net)
    proc_net_remove(net, "dev");
}

-static struct pernet_operations __net_initdata dev_proc_ops = {
+static struct pernet_operations dev_proc_ops = {
    .init = dev_proc_net_init,
    .exit = dev_proc_net_exit,
};

@@ -4325,7 +4325,7 @@ static void __net_exit netdev_exit(struct net *net)
    kfree(net->dev_index_head);
}

-static struct pernet_operations __net_initdata netdev_net_ops = {
+static struct pernet_operations netdev_net_ops = {
    .init = netdev_init,
    .exit = netdev_exit,
};
@@ -4356,7 +4356,7 @@ static void __net_exit default_device_exit(struct net *net)
    rtnl_unlock();
}

-static struct pernet_operations __net_initdata default_device_ops = {
+static struct pernet_operations default_device_ops = {
    .exit = default_device_exit,
};

diff --git a/net/core/dev_mcast.c b/net/core/dev_mcast.c
index 15241cf..ae35405 100644
--- a/net/core/dev_mcast.c
+++ b/net/core/dev_mcast.c
@@ -285,7 +285,7 @@ static void __net_exit dev_mc_net_exit(struct net *net)
    proc_net_remove(net, "dev_mcast");
}

-static struct pernet_operations __net_initdata dev_mc_net_ops = {
+static struct pernet_operations dev_mc_net_ops = {
    .init = dev_mc_net_init,

```

```
.exit = dev_mc_net_exit,  
};  
diff --git a/net/netlink/af_netlink.c b/net/netlink/af_netlink.c  
index 3252729..4f994c0 100644  
--- a/net/netlink/af_netlink.c  
+++ b/net/netlink/af_netlink.c  
@@ -1888,7 +1888,7 @@ static void __net_exit netlink_net_exit(struct net *net)  
#endif  
}  
  
-static struct pernet_operations __net_initdata netlink_net_ops = {  
+static struct pernet_operations netlink_net_ops = {  
.init = netlink_net_init,  
.exit = netlink_net_exit,  
};
```

Subject: Re: [PATCH] remove __net_initdata attribute
Posted by [den](#) on Fri, 26 Oct 2007 11:26:11 GMT

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

Dave, I am sorry, plz disregard the patch :(

- if the file is a part of the kernel, __exit (and pernet_unregister) will never called and should be dropped
- if the file is a part of a module, __init is NOT dropped on the module loading stage

So, we can still make embedded people happy :) Thanks Pavel to pointing me this out.

Denis V. Lunev wrote:

> The discussion with Eric W. Biederman reveals that the concept for
> __net_initdata is bogus. The ->exit method of pernet_operations
> should be called during module exit.
>
> So, discard this attribute.
>
> Signed-off-by: Denis V. Lunev <den@openvz.org>

Subject: Re: [PATCH] remove __net_initdata attribute
Posted by [davem](#) on Fri, 26 Oct 2007 11:34:34 GMT

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

From: "Denis V. Lunev" <den@openvz.org>
Date: Fri, 26 Oct 2007 13:56:52 +0400

> The discussion with Eric W. Biederman reveals that the concept for
> __net_initdata is bogus. The ->exit method of pernet_operations
> should be called during module exit.
>
> So, discard this attribute.
>
> Signed-off-by: Denis V. Lunev <den@openvz.org>

Exactly what I just asked for :-)

Applied, thanks!

Subject: Re: [PATCH] remove __net_initdata attribute
Posted by [davem](#) on Fri, 26 Oct 2007 11:41:13 GMT

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

From: "Denis V. Lunev" <den@sw.ru>
Date: Fri, 26 Oct 2007 15:28:32 +0400

> Dave, I am sorry, plz disregard the patch :(
>
> - if the file is a part of the kernel, __exit (and pernet_unregister)
> will never called and should be dropped
> - if the file is a part of a module, __init is NOT dropped on the
> module loading stage
>
> So, we can still make embedded people happy :) Thanks Pavel to pointing
> me this out.

Ok, I reverted it.
