
Subject: Re: [PATCH RFC 17/31] net: Factor out __dev_alloc_name from dev_alloc_name

Posted by [Benjamin Thery](#) on Mon, 05 Mar 2007 15:29:49 GMT

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

Hello Eric,

See comments about __dev_alloc_name() below.

Regards,
Benjamin

Eric W. Biederman wrote:

```
> From: Eric W. Biederman <ebiederm@xmission.com> - unquoted
>
> When forcibly changing the network namespace of a device
> I need something that can generate a name for the device
> in the new namespace without overwriting the old name.
>
> __dev_alloc_name provides me that functionality.
>
> Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>
> ---
> net/core/dev.c | 44 ++++++-----+
> 1 files changed, 33 insertions(+), 11 deletions(-)
>
> diff --git a/net/core/dev.c b/net/core/dev.c
> index 32fe905..fc0d2af 100644
> --- a/net/core/dev.c
> +++ b/net/core/dev.c
> @@ -655,9 +655,10 @@ int dev_valid_name(const char *name)
> }
>
> /**
> - * dev_alloc_name - allocate a name for a device
> - * @dev: device
> + * __dev_alloc_name - allocate a name for a device
> + * @net: network namespace to allocate the device name in
> + * @name: name format string
> + * @buf: scratch buffer and result name string
> *
> * Passed a format string - eg "l%d" it will try and find a suitable
> * id. It scans list of devices to build up a free map, then chooses
> @@ -668,18 +669,13 @@ int dev_valid_name(const char *name)
> * Returns the number of the unit assigned or a negative errno code.
> */
>
> -int dev_alloc_name(struct net_device *dev, const char *name)
```

```
> +static int __dev_alloc_name(net_t net, const char *name, char buf[IFNAMSIZ])
```

IMHO the third parameter should be: char *buf

Indeed using "char buf[IFNAMSIZ]" is misleading because later in the routine sizeof(buf) is used (with an expected result of IFNAMSIZ).

Unfortunately this is no longer the case: sizeof(buf) value is only 4 now (buf is pointer parameter).

This corrupts the registration of network devices (now I understand why only one of my e1000 showed up after each reboot :).

Also sizeof(buf) should be replaced by IFNAMSIZ in this new routine.
(See below)

```
> {  
>   int i = 0;  
> - char buf[IFNAMSIZ];  
>   const char *p;  
>   const int max_netdevices = 8*PAGE_SIZE;  
>   long *inuse;  
>   struct net_device *d;  
> - net_t net;  
> -  
> - BUG_ON(null_net(dev->nd_net));  
> - net = dev->nd_net;  
>  
>   p = strnchr(name, IFNAMSIZ-1, '%');  
>   if (p) {  
@@ -713,10 +709,8 @@ int dev_alloc_name(struct net_device *dev, const char *name)  
> }  
>  
>   snprintf(buf, sizeof(buf), name, i);
```

Replace "snprintf(buf, IFNAMSIZ, name, i);" or i will never be appended to name and all your ethernet devices will all try to register the name "eth".

There is another occurrence of "snprintf(buf, sizeof(buf), ...)" to replace in the for loop above.

```
> - if (!__dev_get_by_name(net, buf)) {  
> -   strlcpy(dev->name, buf, IFNAMSIZ);  
> + if (!__dev_get_by_name(net, buf))  
>   return i;  
> - }  
>  
> /* It is possible to run out of possible slots  
 >   * when the name is long and there isn't enough space left
```

```

> @@ -725,6 +719,34 @@ int dev_alloc_name(struct net_device *dev, const char *name)
>     return -ENFILE;
> }
>
> /**
> + * dev_alloc_name - allocate a name for a device
> + * @dev: device
> + * @name: name format string
> +
> + * Passed a format string - eg "It%d" it will try and find a suitable
> + * id. It scans list of devices to build up a free map, then chooses
> + * the first empty slot. The caller must hold the dev_base or rtnl lock
> + * while allocating the name and adding the device in order to avoid
> + * duplicates.
> + * Limited to bits_per_byte * page size devices (ie 32K on most platforms).
> + * Returns the number of the unit assigned or a negative errno code.
> + */
> +
> +int dev_alloc_name(struct net_device *dev, const char *name)
> +{
> +    char buf[IFNAMSIZ];
> +    net_t net;
> +    int ret;
> +
> +    BUG_ON(null_net(dev->nd_net));
> +    net = dev->nd_net;
> +    ret = __dev_alloc_name(net, name, buf);
> +    if (ret >= 0)
> +        strcpy(dev->name, buf, IFNAMSIZ);
> +    return ret;
> +}
> +
> +
> /**
> + * dev_change_name - change name of a device

```

--
Benjamin Thery - BULL/DT/Open Software R&D

<http://www.bull.com>

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