
Subject: [PATCH][IPVS] Don't leak sysctl tables if the scheduler registration fails
Posted by [Pavel Emelianov](#) on Mon, 03 Dec 2007 10:04:53 GMT

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

In case we load lbcl or lbclr module we can leak some sysctl tables if the call to register_ip_vs_scheduler() fails.

I've looked at the register_ip_vs_scheduler() code and saw, that the only reason to fail is the name collision, so I think that with some 3rd party schedulers this becomes a relevant issue. No?

Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

```
diff --git a/net/ipv4/ipvs/ip_vs_lbcl.c b/net/ipv4/ipvs/ip_vs_lbcl.c
index b843a11..ad89644 100644
--- a/net/ipv4/ipvs/ip_vs_lbcl.c
+++ b/net/ipv4/ipvs/ip_vs_lbcl.c
@@ -580,9 +580,14 @@ static struct ip_vs_scheduler ip_vs_lbcl_scheduler =
```

```
static int __init ip_vs_lbcl_init(void)
{
+ int ret;
+
  INIT_LIST_HEAD(&ip_vs_lbcl_scheduler.n_list);
  sysctl_header = register_sysctl_table(lbcl_root_table);
- return register_ip_vs_scheduler(&ip_vs_lbcl_scheduler);
+ ret = register_ip_vs_scheduler(&ip_vs_lbcl_scheduler);
+ if (ret)
+ unregister_sysctl_table(sysctl_header);
+ return ret;
}
```

```
diff --git a/net/ipv4/ipvs/ip_vs_lbclr.c b/net/ipv4/ipvs/ip_vs_lbclr.c
index e5b323a..2a5ed85 100644
--- a/net/ipv4/ipvs/ip_vs_lbclr.c
+++ b/net/ipv4/ipvs/ip_vs_lbclr.c
@@ -769,9 +769,14 @@ static struct ip_vs_scheduler ip_vs_lbclr_scheduler =
```

```
static int __init ip_vs_lbclr_init(void)
{
+ int ret;
+
  INIT_LIST_HEAD(&ip_vs_lbclr_scheduler.n_list);
  sysctl_header = register_sysctl_table(lbclr_root_table);
- return register_ip_vs_scheduler(&ip_vs_lbclr_scheduler);
```

```
+ ret = register_ip_vs_scheduler(&ip_vs_lbcr_scheduler);
+ if (ret)
+ unregister_sysctl_table(sysctl_header);
+ return ret;
}
```

Subject: Re: [PATCH][IPVS] Don't leak sysctl tables if the scheduler registration fails

Posted by [Simon Horman](#) on Tue, 04 Dec 2007 01:36:48 GMT

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

On Mon, Dec 03, 2007 at 01:04:53PM +0300, Pavel Emelyanov wrote:

> In case we load lbicr or lbcr module we can leak some sysctl
> tables if the call to register_ip_vs_scheduler() fails.

This looks correct to me.

> I've looked at the register_ip_vs_scheduler() code and saw, that
> the only reason to fail is the name collision, so I think that
> with some 3rd party schedulers this becomes a relevant issue. No?

I guess so. Though presumably register_ip_vs_scheduler() could have other modes of failure in the future.

> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>

Acked-by: Simon Horman <horms@verge.net.au>

```
> ---
>
> diff --git a/net/ipv4/ipvs/ip_vs_lbicr.c b/net/ipv4/ipvs/ip_vs_lbicr.c
> index b843a11..ad89644 100644
> --- a/net/ipv4/ipvs/ip_vs_lbicr.c
> +++ b/net/ipv4/ipvs/ip_vs_lbicr.c
> @@ -580,9 +580,14 @@ static struct ip_vs_scheduler ip_vs_lbicr_scheduler =
>
> static int __init ip_vs_lbicr_init(void)
> {
> + int ret;
> +
> + INIT_LIST_HEAD(&ip_vs_lbicr_scheduler.n_list);
> + sysctl_header = register_sysctl_table(lbicr_root_table);
> - return register_ip_vs_scheduler(&ip_vs_lbicr_scheduler);
> + ret = register_ip_vs_scheduler(&ip_vs_lbicr_scheduler);
> + if (ret)
> + unregister_sysctl_table(sysctl_header);
> + return ret;
```

```
> }
>
>
> diff --git a/net/ipv4/ipvs/ip_vs_lblcr.c b/net/ipv4/ipvs/ip_vs_lblcr.c
> index e5b323a..2a5ed85 100644
> --- a/net/ipv4/ipvs/ip_vs_lblcr.c
> +++ b/net/ipv4/ipvs/ip_vs_lblcr.c
> @@ -769,9 +769,14 @@ static struct ip_vs_scheduler ip_vs_lblcr_scheduler =
>
> static int __init ip_vs_lblcr_init(void)
> {
> + int ret;
> +
> + INIT_LIST_HEAD(&ip_vs_lblcr_scheduler.n_list);
> + sysctl_header = register_sysctl_table(lblcr_root_table);
> - return register_ip_vs_scheduler(&ip_vs_lblcr_scheduler);
> + ret = register_ip_vs_scheduler(&ip_vs_lblcr_scheduler);
> + if (ret)
> + unregister_sysctl_table(sysctl_header);
> + return ret;
> }
>
>
```

--
Horms

Subject: Re: [PATCH][IPVS] Don't leak sysctl tables if the scheduler registration fails

Posted by [davem](#) on Tue, 04 Dec 2007 08:43:50 GMT

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

From: Simon Horman <horms@verge.net.au>

Date: Tue, 4 Dec 2007 10:36:48 +0900

> On Mon, Dec 03, 2007 at 01:04:53PM +0300, Pavel Emelyanov wrote:

> > In case we load lblc or lblcr module we can leak some sysctl
> > tables if the call to register_ip_vs_scheduler() fails.

>
> This looks correct to me.

>
> > I've looked at the register_ip_vs_scheduler() code and saw, that
> > the only reason to fail is the name collision, so I think that
> > with some 3rd party schedulers this becomes a relevant issue. No?

>
> I guess so. Though presumably register_ip_vs_scheduler() could
> have other modes of failure in the future.

>
> > Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
>
> Acked-by: Simon Horman <horms@verge.net.au>

Applied, thanks.
