
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_lbr_scheduler);  
+ if (ret)  
+ unregister_sysctl_table(sysctl_header);  
+ return ret;  
}
```
