

---

Subject: [PATCH 1/2][INET] Fix potential kfree on vmalloc-ed area of request\_sock\_queue

Posted by Pavel Emelianov on Wed, 14 Nov 2007 18:08:29 GMT

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

---

The request\_sock\_queue's listen\_opt is either vmalloc-ed or kmalloc-ed depending on the number of table entries. Thus it is expected to be handled properly on free, which is done in the reqsk\_queue\_destroy().

However the error path in inet\_csk\_listen\_start() calls the lite version of reqsk\_queue\_destroy, called \_\_reqsk\_queue\_destroy, which calls the kfree unconditionally.

Fix this and move the \_\_reqsk\_queue\_destroy into a .c file as it looks too big to be inline.

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

---

```
diff --git a/include/net/request_sock.h b/include/net/request_sock.h
index 7aed02c..0a954ee 100644
--- a/include/net/request_sock.h
+++ b/include/net/request_sock.h
@@ -136,11 +136,7 @@ static inline struct listen_sock *reqsk_queue_yank_listen_sk(struct
request_sock
    return lopt;
}
```

```
-static inline void __reqsk_queue_destroy(struct request_sock_queue *queue)
-{
-    kfree(reqsk_queue_yank_listen_sk(queue));
-}
-
+extern void __reqsk_queue_destroy(struct request_sock_queue *queue);
extern void reqsk_queue_destroy(struct request_sock_queue *queue);
```

```
static inline struct request_sock *
diff --git a/net/core/request_sock.c b/net/core/request_sock.c
index 5f0818d..12a15f5 100644
--- a/net/core/request_sock.c
+++ b/net/core/request_sock.c
@@ -71,6 +71,20 @@ int reqsk_queue_alloc(struct request_sock_queue *queue,
```

EXPORT\_SYMBOL(reqsk\_queue\_alloc);

```
+void __reqsk_queue_destroy(struct request_sock_queue *queue)
```

```
+{
+ struct listen_sock *lopt = reqsk_queue_yank_listen_sk(queue);
+ size_t lopt_size = sizeof(struct listen_sock) +
+ lopt->nr_table_entries * sizeof(struct request_sock *);
+
+ if (lopt_size > PAGE_SIZE)
+ vfree(lopt);
+ else
+ kfree(lopt);
+}
+
+EXPORT_SYMBOL(__reqsk_queue_destroy);
+
void reqsk_queue_destroy(struct request_sock_queue *queue)
{
/* make all the listen_opt local to us */
```

---