Subject: [PATCH netdev] "wrong timeout value" in sk_wait_data()
Posted by vaverin on Thu, 24 May 2007 04:22:03 GMT
View Forum Message <> Reply to Message

sys_setsockopt() do not check properly timeout values for
SO_RCVTIMEO/SO_SNDTIMEO, for example it's possible to set negative timeout
values. POSIX do not defines behaviour for sys_setsockopt in case negative
timeouts, but requires that setsockopt() shall fail with -EDOM if the send and
receive timeout values are too big to fit into the timeout fields in the socket
structure.
In current implementation negative timeout can lead to error messages like
"schedule_timeout: wrong timeout value".

Proposed patch:
- checks tv_usec and returns -EDOM if it is wrong
- do not allows to set negative timeout values (sets 0 instead) and outputs
ratelimited information message about such attempts.

Signed-Off-By: Vasily Averin <vvs@sw.ru>

```
diff --git a/net/core/sock.c b/net/core/sock.c
index 22183c2..27d7a46 100644
--- a/net/core/sock.c
+++ b/net/core/sock.c
@@ -206,7 +206,19 @@ static int sock_set_timeout(long *timeo_p, char __user
*optval, int optlen)
   return -EINVAL;
  if (copy_from_user(&tv, optval, sizeof(tv)))
   return -EFAULT;
-
+ if (tv.tv_usec < 0 || tv.tv_usec >= 1000000)
+  return -EDOM;
+
+ if (tv.tv_sec < 0) {
+  static int warned = 0;
+  *timeo_p = 0;
+  if (warned < 10 && net_ratelimit())
+   warned++;
+   printk(KERN_INFO "sock_set_timeout: `%s' (pid %d) "
+          "tries to set negative timeout\n",
+           current->comm, current->pid);
+  return 0;
+ }
  *timeo_p = MAX_SCHEDULE_TIMEOUT;
  if (tv.tv_sec == 0 && tv.tv_usec == 0)
   return 0;
```

Subject: Re: [PATCH netdev] "wrong timeout value" in sk_wait_data()
Posted by Eric Dumazet on Thu, 24 May 2007 04:36:05 GMT
View Forum Message <> Reply to Message

Vasily Averin a écrit :
> sys_setsockopt() do not check properly timeout values for
> SO_RCVTIMEO/SO_SNDTIMEO, for example it's possible to set negative timeout
> values. POSIX do not defines behaviour for sys_setsockopt in case negative
> timeouts, but requires that setsockopt() shall fail with -EDOM if the send and
> receive timeout values are too big to fit into the timeout fields in the socket
> structure.
> In current implementation negative timeout can lead to error messages like
> "schedule_timeout: wrong timeout value".
>
> Proposed patch:
> - checks tv_usec and returns -EDOM if it is wrong
> - do not allows to set negative timeout values (sets 0 instead) and outputs
> ratelimited information message about such attempts.
>
> Signed-Off-By: Vasily Averin <vvs@sw.ru>
>
> diff --git a/net/core/sock.c b/net/core/sock.c
> index 22183c2..27d7a46 100644
> --- a/net/core/sock.c
> +++ b/net/core/sock.c
> @@ -206,7 +206,19 @@ static int sock_set_timeout(long *timeo_p, char __user
> *optval, int optlen)
> return -EINVAL;
> if (copy_from_user(&tv, optval, sizeof(tv)))
> return -EFAULT;
> -
> + if (tv.tv_usec < 0 || tv.tv_usec >= 1000000)
> + return -EDOM;

Please use USEC_PER_SEC instead of 1000000

> +
> + if (tv.tv_sec < 0) {
> + static int warned = 0;
> + *timeo_p = 0;
> + if (warned < 10 && net_ratelimit())
> + warned++;
> + printk(KERN_INFO "sock_set_timeout: `%s' (pid %d) "
> + "tries to set negative timeout\n",
> + current->comm, current->pid);
> + return 0;
> + }
> *timeo_p = MAX_SCHEDULE_TIMEOUT;
> if (tv.tv_sec == 0 && tv.tv_usec == 0)

> return 0;
>

Subject: [PATCH netdev] "wrong timeout value" in sk_wait_data() v2
Posted by vaverin on Thu, 24 May 2007 05:23:14 GMT
View Forum Message <> Reply to Message

sys_setsockopt() do not check properly timeout values for
SO_RCVTIMEO/SO_SNDTIMEO, for example it's possible to set negative timeout
values. POSIX do not defines behaviour for sys_setsockopt in case negative
timeouts, but requires that setsockopt() shall fail with -EDOM if the send and
receive timeout values are too big to fit into the timeout fields in the socket
structure.
In current implementation negative timeout can lead to error messages like
"schedule_timeout: wrong timeout value".

Proposed patch:
- checks tv_usec and returns -EDOM if it is wrong
- do not allows to set negative timeout values (sets 0 instead) and outputs
ratelimited information message about such attempts.

Signed-Off-By: Vasily Averin <vvs@sw.ru>

```
diff --git a/net/core/sock.c b/net/core/sock.c
index 22183c2..7e51d3a 100644
--- a/net/core/sock.c
+++ b/net/core/sock.c
@@ -206,7 +206,19 @@ static int sock_set_timeout(long *timeo_p, char __user
*optval, int optlen)
   return -EINVAL;
  if (copy_from_user(&tv, optval, sizeof(tv)))
   return -EFAULT;
-
+ if (tv.tv_usec < 0 || tv.tv_usec >= USEC_PER_SEC)
+  return -EDOM;
+
+ if (tv.tv_sec < 0) {
+  static int warned = 0;
+  *timeo_p = 0;
+  if (warned < 10 && net_ratelimit())
+   warned++;
+   printk(KERN_INFO "sock_set_timeout: `%s' (pid %d) "
+          "tries to set negative timeout\n",
+          current->comm, current->pid);
+  return 0;
+ }
  *timeo_p = MAX_SCHEDULE_TIMEOUT;
```

```
 if (tv.tv_sec == 0 && tv.tv_usec == 0)
  return 0;
```

Subject: Re: [PATCH netdev] "wrong timeout value" in sk_wait_data() v2
Posted by davem on Thu, 24 May 2007 23:59:47 GMT
View Forum Message <> Reply to Message

From: Vasily Averin <vvs@sw.ru>
Date: Thu, 24 May 2007 09:23:14 +0400

> sys_setsockopt() do not check properly timeout values for
> SO_RCVTIMEO/SO_SNDTIMEO, for example it's possible to set negative timeout
> values. POSIX do not defines behaviour for sys_setsockopt in case negative
> timeouts, but requires that setsockopt() shall fail with -EDOM if the send and
> receive timeout values are too big to fit into the timeout fields in the socket
> structure.
> In current implementation negative timeout can lead to error messages like
> "schedule_timeout: wrong timeout value".
>
> Proposed patch:
> - checks tv_usec and returns -EDOM if it is wrong
> - do not allows to set negative timeout values (sets 0 instead) and outputs
> ratelimited information message about such attempts.
>
> Signed-Off-By: Vasily Averin <vvs@sw.ru>

Thank you for this bug fix, patch applied.