

Re: [PATCH] Fix user data corrupted by old value return of sysctl

# Re: [PATCH] Fix user data corrupted by old value return of sysctl

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-12/msg08538.html>

---

- *From:* Yi Yang <[yang.y.yi@xxxxxxxxxx](mailto:yang.y.yi@xxxxxxxxxx)>
  - *Date:* Sat, 31 Dec 2005 09:08:54 +0800
- 

Linus Torvalds wrote:

On Fri, 30 Dec 2005, Yi Yang wrote:

```
If the user reads a sysctl entry which is of string
type
by sysctl syscall, this call probably corrupts the
user data
right after the old value buffer, the issue lies in
sysctl_string
setting 0 to oldval[len], len is the available
buffer size
specified by the user, obviously, this will write
to the first
byte of the user memory place immediate after the
old value buffer,
the correct way is that sysctl_string doesn't set
0, the user
should do it by self in the program.
```

Hmm.. I think this patch is incomplete.

We `_should_` zero-pad the data, at least if the result fits in the buffer.

So I think the correct fix is to just `_copy_` the last zero if it fits in the buffer, rather than do the unconditional "add NUL at the end" thing. The simplest way to do that is to just make "l" be "strlen(str)+1", so that we count the ending NUL

Re: [PATCH] Fix user data corrupted by old value return of sysctl

in the length (and then, if the buffer isn't big enough, we will truncate it).

In other words, I would instead suggest a patch like the appended.

But even that is questionable: one alternative is to always zero-pad (like we used to), but make sure that the buffer size is sufficient for it (ie instead of adding one to the length of the string, we'd subtract one from the buffer length and make sure that the '\0' fits..

Comments?

Yes, you are more complete, I agree with it very much.

Linus

```
---
diff --git a/kernel/sysctl.c b/kernel/sysctl.c
index 9990e10..ad0425a 100644
--- a/kernel/sysctl.c
+++ b/kernel/sysctl.c
@@ -2201,14 +2201,12 @@ int sysctl_string(ctl_table *table,
int if (get_user(len, oldlenp))
return -EFAULT;
if (len) {
- l = strlen(table->data);
+ l = strlen(table->data)+1;
if (len > l) len = l;
if (len >= table->maxlen)
len = table->maxlen;
if(copy_to_user(oldval, table->data, len))
return -EFAULT;
- if(put_user(0, ((char __user *) oldval) + len))
- return -EFAULT;
if(put_user(len, oldlenp))
return -EFAULT;
}
```

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxxxxx  
More majordomo info at <http://vqer.kernel.org/majordomo-info.html>

Re: [PATCH] Fix user data corrupted by old value return of sysctl

Please read the FAQ at <http://www.tux.org/lkml/>