
Subject: [PATCH 2/3] sysctl: Factor out sysctl_data.
Posted by [ebiederm](#) on Thu, 09 Aug 2007 19:52:16 GMT
[View Forum Message](#) <> [Reply to Message](#)

There as been no easy way to wrap the default sysctl strategy routine except for returning 0. Which is not always what we want. The few instances I have seen that want different behaviour have written their own version of sysctl_data. While not too hard it is unnecessary code and has the potential for extra bugs.

So to make these situations easier and make that part of sysctl more symmetric I have factored sysctl_data out of do_sysctl_strategy and exported as a function everyone can use.

Further having sysctl_data be an explicit function makes checking for badly formed sysctl tables much easier.

Signed-off-by: Eric W. Biederman <ebiederm@xmission.com>

```
include/linux/sysctl.h |  1 +
kernel/sysctl.c      | 66 ++++++-----+
2 files changed, 47 insertions(+), 20 deletions(-)
```

```
diff --git a/include/linux/sysctl.h b/include/linux/sysctl.h
index f73be4c..5ca510b 100644
--- a/include/linux/sysctl.h
+++ b/include/linux/sysctl.h
@@ -972,6 +972,7 @@ extern int do_sysctl_strategy (struct ctl_table *table,
    void __user *oldval, size_t __user *oldlenp,
    void __user *newval, size_t newlen);
```

```
+extern ctl_handler sysctl_data;
extern ctl_handler sysctl_string;
extern ctl_handler sysctl_intvec;
extern ctl_handler sysctl_jiffies;
diff --git a/kernel/sysctl.c b/kernel/sysctl.c
index 6723f92..cf4c632 100644
--- a/kernel/sysctl.c
+++ b/kernel/sysctl.c
@@ -1406,7 +1406,6 @@ int do_sysctl_strategy (struct ctl_table *table,
    void __user *newval, size_t newlen)
{
    int op = 0, rc;
-   size_t len;

    if (oldval)
        op |= 004;
@@ -1427,25 +1426,10 @@ int do_sysctl_strategy (struct ctl_table *table,
```

```

/* If there is no strategy routine, or if the strategy returns
 * zero, proceed with automatic r/w */
if (table->data && table->maxlen) {
- if (oldval && oldlenp) {
-   if (get_user(len, oldlenp))
-     return -EFAULT;
-   if (len) {
-     if (len > table->maxlen)
-       len = table->maxlen;
-     if(copy_to_user(oldval, table->data, len))
-       return -EFAULT;
-     if(put_user(len, oldlenp))
-       return -EFAULT;
-   }
- }
- if (newval && newlen) {
-   len = newlen;
-   if (len > table->maxlen)
-     len = table->maxlen;
-   if(copy_from_user(table->data, newval, len))
-     return -EFAULT;
- }
+ rc = sysctl_data(table, name, nlen, oldval, oldlenp,
+   newval, newlen);
+ if (rc < 0)
+   return rc;
}
return 0;
}
@@ -2344,6 +2328,40 @@ int proc_doulongvec_ms_jiffies_minmax(struct ctl_table *table, int
write,
 * General sysctl support routines
 */
+/* The generic sysctl data routine (used if no strategy routine supplied) */
+int sysctl_data(struct ctl_table *table, int __user *name, int nlen,
+  void __user *oldval, size_t __user *oldlenp,
+  void __user *newval, size_t newlen)
+{
+  size_t len;
+
+  /* Get out of I don't have a variable */
+  if (!table->data || !table->maxlen)
+    return -ENOTDIR;
+
+  if (oldval && oldlenp) {
+    if (get_user(len, oldlenp))
+      return -EFAULT;

```

```

+ if (len) {
+   if (len > table->maxlen)
+     len = table->maxlen;
+   if (copy_to_user(oldval, table->data, len))
+     return -EFAULT;
+   if (put_user(len, oldlenp))
+     return -EFAULT;
+ }
+ }
+
+ if (newval && newlen) {
+   if (newlen > table->maxlen)
+     newlen = table->maxlen;
+   if (copy_from_user(table->data, newval, newlen))
+     return -EFAULT;
+ }
+ return 1;
+}
+
/* The generic string strategy routine: */
int sysctl_string(struct ctl_table *table, int __user *name, int nlen,
    void __user *oldval, size_t __user *oldlenp,
@@ -2532,6 +2550,13 @@ out:
    return -ENOSYS;
}

+int sysctl_data(struct ctl_table *table, int __user *name, int nlen,
+    void __user *oldval, size_t __user *oldlenp,
+    void __user *newval, size_t newlen)
+{
+    return -ENOSYS;
+}
+
int sysctl_string(struct ctl_table *table, int __user *name, int nlen,
    void __user *oldval, size_t __user *oldlenp,
    void __user *newval, size_t newlen)
@@ -2579,4 +2604,5 @@ EXPORT_SYMBOL(sysctl_intvec);
EXPORT_SYMBOL(sysctl_jiffies);
EXPORT_SYMBOL(sysctl_ms_jiffies);
EXPORT_SYMBOL(sysctl_string);
+EXPORT_SYMBOL(sysctl_data);
EXPORT_SYMBOL(unregister_sysctl_table);
--
```

1.5.1.1.181.g2de0
