

## Re: [git pull] kgdb light, v5

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-02/msg04837.html>

---

- *From:* Ingo Molnar <[mingo@xxxxxxx](mailto:mingo@xxxxxxx)>
  - *Date:* Sun, 10 Feb 2008 21:29:30 +0100
- 

\* Linus Torvalds <[torvalds@xxxxxxxxxxxxxxxxxxxxxxxx](mailto:torvalds@xxxxxxxxxxxxxxxxxxxxxxxx)> wrote:

```
+static int kgdb_get_mem(char *addr, unsigned char *buf, int count)
{
+ if ((unsigned long)addr < TASK_SIZE)
+ return -EFAULT;

+ return probe_kernel_read(buf, addr, count);
}
```

Ok, so this is a pretty function after all the cleanups, but I actually don't think that "if ((unsigned long)addr < TASK\_SIZE)" is really even asked for.

Why not let kgdb look at user memory? I'd argue that in a lot of cases, it might be quite nice to do, to see what user arguments in memory are etc etc (think things like futexes, where user memory contents really do matter).

So I'd suggest getting rid of the whole "kgdb\_{get|set}\_mem()" functions, and just using "probe\_kernel\_{read|write}()" directly instead.

ok, on a second thought: kgdb\_{get|set}\_mem() is *\_only\_* used to validate and set the software breakpoint (int3). And i think kgdb correctly restricts that to kernel-space addresses only – you can typo an address down into user-space and overwrite user-space memory and not know what hit you ... [you can still explicitly touch user-space memory, but that has to be done intentionally]

So to reduce the confusion i've removed these functions and open-coded the probe\_kernel\_\*(\*) functions into kgdb\_validate\_break\_address() and kgdb\_arch\_set\_breakpoint().

all other places already use probe\_kernel\_{read|write}. (Now, there are

Re: [git pull] kgdb light, v5

a few stray TASK\_SIZE checks still, i'll double check them and convert them to access\_ok() checks.)

btw., based on your previous comment about alignment, i found another function that used weird alignment checks, kgdb\_hex2mem():

```
if (count == 2 && ((long)mem & 1) == 0)
err = probe_kernel_write(mem, tmp_raw, 2);
else if (count == 4 && ((long)mem & 3) == 0)
err = probe_kernel_write(mem, tmp_raw, 4);
else if (count == 8 && ((long)mem & 7) == 0)
err = probe_kernel_write(mem, tmp_raw, 8);
else
err = probe_kernel_write(mem, tmp_raw, count);

return err;
}
```

I just converted it to:

```
return probe_kernel_write(mem, tmp_raw, count);
```

which looks \_a lot\_ cleaner.

Ingo

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

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