

# Re: [PATCH] Use BIOS reboot on Toshiba Portege 4000

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-11/msg00911.html>

---

- *From:* [ebiederm@xxxxxxxxxxxxx](mailto:ebiederm@xxxxxxxxxxxxx) (Eric W. Biederman)
  - *Date:* Tue, 04 Nov 2008 03:22:30 -0800
- 

Avi Kivity <[avi@xxxxxxxxxxxxx](mailto:avi@xxxxxxxxxxxxx)> writes:

Eric W. Biederman wrote:

I think we are confusing two issues here.

– Ordinary `machine_restart` which happens to call `emergency_restart`.  
And is proceeded by `machine_shutdown`.

– And `emergency_restart` itself.

To some extent I would be a lot happier if `Alt-sysrq-r` did what was necessary to get into a context where it can call `machine_restart` or even better `kernel_restart()`.  
`emergency_restart()` is a nice idea but is broken by design.

Isn't `emergency_restart()` equivalent to `kexec()`? Both start from indeterminate conditions.

Good point. That is a reasonable direction to evolve it on x86.  
Similar to and sharing some of the same code as the `kexec` on panic path.

We may need to separate out `emergency_restart` from the normal clean restart to make that happen. It would be pointless and silly to be sending NMI at other cpus for example if we have cleanly shut them down already.

That said. If we can turn off `vmx` on that one processor.  
That should be enough for the cpu to triple fault and let the BIOS do what it needs to do on that cpu i.e. `outb(magic, 0x92)` and toggle a motherboard level reset?

Re: [PATCH] Use BIOS reboot on Toshiba Portege 4000

If triple fault is wired to INIT (as it is at least on some systems; for example one of mine) then the cpu will reset, but why will the bios proceed to issue a motherboard reset? Won't it happily POST it's way to boot (leaving the other cpus dead)?

I'm not certain. But when I was writing BIOS's it was much easier to just toggle the reset line than to try and cope with the weird state the machine was in. I'm pretty certain why we don't see more problems with reboot when we leave the machine in a weird state. It is certainly legal for a BIOS to just run the POST code though.

If I read the earlier comments correctly the deep issue is that going through ACPI to reset systems is less reliable than doing it the classic way.

It depends on the system; both are unreliable. But if we use the same trick as with kdump (NMI SIPI + vmxoff) the choice will be orthogonal to whether vmx is in use or not.

Yes.

Eric

--

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

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

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