

Solving suspend-level confusion

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-07/6202.html>

From: Pavel Machek (pavel_at_suse.cz)

Date: 07/30/04

Date: Fri, 30 Jul 2004 18:44:13 +0200

To: benh@kernel.crashing.org, kernel list <linux-kernel@vger.kernel.org>, Patrick Mochel <[Hi!](mailto:mochel@></p></div><div data-bbox=)

There's quite big confusion what 'u32 state' in suspend callbacks mean. Half code interprets it as a system-wide suspend level, and second half thinks it is PCI Dx state. Bad.

What about this solution:

- * system-wide suspend level is always passed down (it is more detailed, and for example IDE driver cares)

- * if you want to derive Dx state, just use provided inline function to convert.

That should be acceptable to everyone. I plan on explicitly using enums to prevent further confusion. Patch is likely to be pretty big: ideally all prototypes of suspend function would be fixed so noone can be confused.

What do you think?

Pavel

PS: I'll be on holidays for a week, feel free to implement this or something similar.. it is going to be lot of search/replace in drivers :-(.

```
--- tmp/linux/include/linux/pci.h 2004-06-22 12:36:45.000000000 +0200
```

```
+++ linux/include/linux/pci.h 2004-07-30 18:24:02.000000000 +0200
```

```
@@ -593,7 +593,7 @@
```

```
    const struct pci_device_id *id_table; /* must be non-NULL for probe to be called */
    int (*probe) (struct pci_dev *dev, const struct pci_device_id *id); /* New device inserted */
    void (*remove) (struct pci_dev *dev); /* Device removed (NULL if not a hot-plug capable driver) */
- int (*suspend) (struct pci_dev *dev, u32 state); /* Device suspended */
+ int (*suspend) (struct pci_dev *dev, enum suspend_state reason); /* Device suspended */
    int (*resume) (struct pci_dev *dev); /* Device woken up */
    int (*enable_wake) (struct pci_dev *dev, u32 state, int enable); /* Enable wake event */
```

```
--- tmp/linux/include/linux/pm.h 2004-06-22 12:36:45.000000000 +0200
```

Linux-Kernel: Solving suspend-level confusion

```
+++ linux/include/linux/pm.h 2004-07-30 18:23:11.000000000 +0200
@@ -193,11 +193,11 @@
extern void (*pm_idle)(void);
extern void (*pm_power_off)(void);

-enum {
- PM_SUSPEND_ON,
- PM_SUSPEND_STANDBY,
- PM_SUSPEND_MEM,
- PM_SUSPEND_DISK,
+enum system_state {
+ PM_SUSPEND_ON = 0,
+ PM_SUSPEND_STANDBY = 1,
+ PM_SUSPEND_MEM = 2,
+ PM_SUSPEND_DISK = 3,
+     PM_SUSPEND_MAX,
};

@@ -241,11 +240,47 @@

extern void device_pm_set_parent(struct device * dev, struct device * parent);

-extern int device_suspend(u32 state);
-extern int device_power_down(u32 state);
+enum suspend_state {
+ SNAPSHOT = 10, /* For debugging, symbolic constants should be everywhere */
+ POWERDOWN,
+ RESTART,
+ RUNTIME_D1,
+ RUNTIME_D2,
+ RUNTIME_D3hot,
+ RUNTIME_D3cold
+};
+
+extern int device_suspend(enum suspend_state reason);
+extern int device_power_down(enum suspend_state reason);
extern void device_power_up(void);
extern void device_resume(void);

+enum pci_state {
+ D0 = 20, /* For debugging, symbolic constants should be everywhere */
+ D1,
+ D2,
+ D3hot,
+ D3cold
+};
+
+static inline enum pci_state to_pci_state(enum suspend_state state)
+{
+ switch(state) {
+ case RUNTIME_D1:
```

Linux-Kernel: Solving suspend-level confusion

```
+ return D1;
+ case RUNTIME_D2:
+ return D2;
+ case RUNTIME_D3hot:
+ return D3hot;
+ case SNAPSHOT:
+ case POWERDOWN:
+ case RESTART:
+ case RUNTIME_D3cold:
+ return D3cold;
+ default:
+ BUG();
+ }
+ }
```

```
#endif /* __KERNEL__ */
```

```
--- tmp/linux/kernel/power/disk.c 2004-05-20 23:08:36.000000000 +0200
```

```
+++ linux/kernel/power/disk.c 2004-07-30 18:18:04.000000000 +0200
```

```
@@ -46,20 +46,25 @@
```

```
int error = 0;
```

```
local_irq_save(flags);
```

```
- device_power_down(PM_SUSPEND_DISK);
```

```
switch(mode) {
```

```
case PM_DISK_PLATFORM:
```

```
+ device_power_down(POWERDOWN);
```

```
error = pm_ops->enter(PM_SUSPEND_DISK);
```

```
break;
```

```
case PM_DISK_SHUTDOWN:
```

```
+ device_power_down(POWERDOWN);
```

```
printk("Powering off system\n");
```

```
machine_power_off();
```

```
break;
```

```
case PM_DISK_REBOOT:
```

```
+ device_power_down(RESTART);
```

```
machine_restart(NULL);
```

```
break;
```

```
}
```

```
machine_halt();
```

```
+ /* Valid image is on the disk, if we continue we risk serious data corruption
```

```
+ after resume. */
```

```
+ while(1);
```

```
device_power_up();
```

```
local_irq_restore(flags);
```

```
return 0;
```

```
--
```

```
People were complaining that M$ turns users into beta-testers...
```

```
...jr ghea gurz vagb qrirybcref, naq gurl frrz gb yvxr vg gung jnl!
```

```
-
```

```
To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
```

Linux–Kernel: Solving suspend–level confusion

the body of a message to majordomo@vger.kernel.org

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

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