Fixing suspend for fun and profit

(But not necessarily all that much profit)

Matthew Garrett
<mjg59@ubuntu.com>

ubuntu

Fruitflies

ubuntu

Kernels

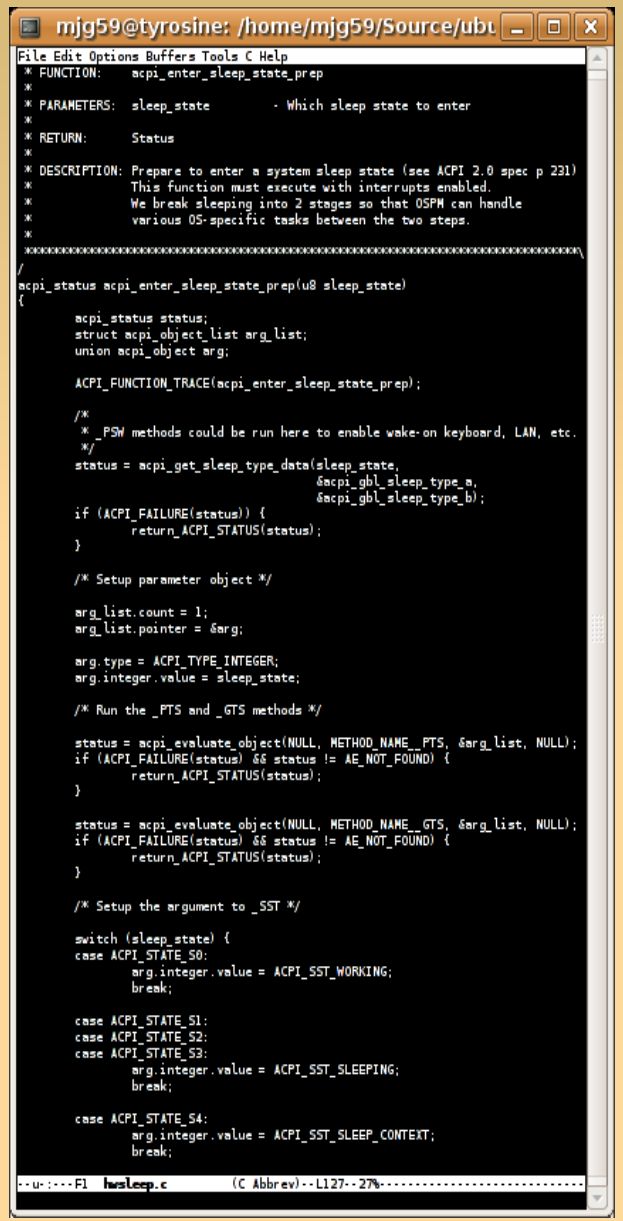


```
 * FUNCTION:    acpi_enter_sleep_state_prep
 *
 * PARAMETERS:  sleep_state         - Which sleep state to enter
 *
 * RETURN:      Status
 *
 * DESCRIPTION: Prepare to enter a system sleep state (see ACPI 2.0 spec p 231)
 *              This function must execute with interrupts enabled.
 *              We break sleeping into 2 stages so that OSPM can handle
 *              various OS-specific tasks between the two steps.
 *
 ******************************************************************************\
/
acpi_status acpi_enter_sleep_state_prep(u8 sleep_state)
{
        acpi_status status;
        struct acpi_object_list arg_list;
        union acpi_object arg;

        ACPI_FUNCTION_TRACE(acpi_enter_sleep_state_prep);

        /*
         * _PSW methods could be run here to enable wake-on keyboard, LAN, etc.
         */
        status = acpi_get_sleep_type_data(sleep_state,
                                          &acpi_gbl_sleep_type_a,
                                          &acpi_gbl_sleep_type_b);
        if (ACPI_FAILURE(status)) {
                return_ACPI_STATUS(status);
        }

        /* Setup parameter object */

        arg_list.count = 1;
        arg_list.pointer = &arg;

        arg.type = ACPI_TYPE_INTEGER;
        arg.integer.value = sleep_state;

        /* Run the _PTS and _GTS methods */

        status = acpi_evaluate_object(NULL, METHOD_NAME__PTS, &arg_list, NULL);
        if (ACPI_FAILURE(status) && status != AE_NOT_FOUND) {
                return_ACPI_STATUS(status);
        }

        status = acpi_evaluate_object(NULL, METHOD_NAME__GTS, &arg_list, NULL);
        if (ACPI_FAILURE(status) && status != AE_NOT_FOUND) {
                return_ACPI_STATUS(status);
        }

        /* Setup the argument to _SST */

        switch (sleep_state) {
        case ACPI_STATE_S0:
                arg.integer.value = ACPI_SST_WORKING;
                break;

        case ACPI_STATE_S1:
        case ACPI_STATE_S2:
        case ACPI_STATE_S3:
                arg.integer.value = ACPI_SST_SLEEPING;
                break;

        case ACPI_STATE_S4:
                arg.integer.value = ACPI_SST_SLEEP_CONTEXT;
                break;
```

ubuntu

Life

ubuntu

Laptops

ubuntu

Why does suspend not work?

- Drivers
- Drivers
- Kernel
- Drivers

What happens over suspend?

- Stop userspace
- Suspend drivers
- Time stops
- Resume drivers
- Start userspace

What does a driver have to do?

- Quiesce hardware
- Save state
- Put it into an appropriate power state

What does a driver have to do?

- Power up the device
- Restore state

So why is this hard?

After all, it's just restoring saved state

ubuntu

So why is this hard?

After all, it's “just” restoring saved state

ubuntu

A typical PCI driver

- pci_save_state()
- pci_set_power_state()
- pci_restore_state()

pci_save_state

Does **not** save all of the PCI configuration registers

ubuntu

Why not?

Order matters

How can I test my driver?

echo -n 2
>/sys/bus/pci/devices/foo/power/state

(Except somebody broke it)

But my driver works for me!

Not all platforms will leave the hardware in the same state

And yes, it's because they hate you/you do bad things/kittens

Not every device is PCI

Platform drivers face much the same issue

(APIC register state, for example)

ubuntu

MY LAPTOP DOESN'T WORK.

Sorry.

No, really, I'm sorry.

ubuntu

But...

Oh, you'd like it to work?

ubuntu

Blank screen debugging

- CONFIG_PM_TRACE
- TRACE_DEVICE
- TRACE_RESUME
- Utterly, utterly magic

CONFIG_PM_TRACE

echo 1 > /sys/power/pm_trace

Magic number: 4:156:725
hash matches drivers/base/power/resume.c:28
hash matches device 0000:01:00.0

ubuntu

PM is hard. Let's go shopping

Suspend to disk generally works better

(Unfortunately)

ubuntu

This isn't impossible

Even if some vendors try to make it as hard as they can

(Thanks, ATI)

ubuntu