

Power Management on Linux (aka Another Fine Mess)

Sitsofe Wheeler

In the beginning (the x86 perspective)

- ...there was APM
 - BIOS did everything
 - Not very configurable (what your BIOS didn't offer you couldn't do)
 - Pushing the power button powered off your machine

The big change

- Now there's ACPI
 - More operating system control
 - Choose to power off bits of the system to save power
 - Suspend to disk can be made reliable
 - Sends events to the OS
 - No need for special wake up cables
 - Spec is **huge** (over 600 pages)...

[The ACPI spec is a] classical spec in the sense that someone was bribed to go to Honolulu to "talk the spec over" and "reach a compromise". [... My] engineering assessment is ACPI is a pile of camel pooh.

Marco Peereboom (OpenBSD ACPI developer)

The ACPI aftermath

- Windows
 - Everyone tests under Windows...
 - ...and stops there

The ACPI aftermath

- Windows
 - Everyone tests under Windows...
 - ...and stops there
 - Many BIOSes check for the presence of Windows NT/XP before fixing or enabling certain behaviour.

The ACPI aftermath

- Windows
 - Everyone tests under Windows...
 - ...and stops there
 - Many BIOSes check for the presence of Windows NT/XP before fixing or enabling certain behaviour
 - So everyone has to pretend to be Windows NT/XP

The ACPI aftermath

- Windows
 - Everyone tests under Windows...
 - ...and stops there
 - Many BIOSes check for the presence of Windows NT/XP before fixing or enabling certain behaviour
 - So everyone has to pretend to be Windows NT/XP
 - **Even Vista**

[The flaw...] became obvious when Microsoft changed the string returned by _OS to something other than Microsoft Windows NT. [...] Changing to an unknown return string caused adverse effects on these systems.

Microsoft ACPI Documentation

The bother with BIOSes...

- Buggy BIOSes are rife (especially in desktops)
- Different hardware needs different workarounds to get it going
- BIOS ACPI code frequently compiles with lots of warnings
- If you're lucky there will be a BIOS update fixing problems
- Or maybe the update will just break something else...

“Avoid the the A11 revision of the bios on this machine as it stops the USB keyboard working with grub!”

Sitsofe Wheeler on the S2ram mailing list

What usually works

- Booting usually works (at least past ACPI)
 - Some hardware (especially multi CPU computers) now needs ACPI to even boot
- Putting the monitor to sleep
- Pushing the power button sends an event
- Err...

What often works

- Suspend to disk (swsusp, uswsusp, suspend2)
 - Became much more reliable around 2.6.16
 - Works best with driver support but can sometimes be worked around by unloading modules
 - Binary drivers might need extra care
- CPU throttling/scaling
 - Used to have only have idle or in use
 - Then we had idle, 50%, 100%
 - Now we have CPU scaling. Speed can chosen based on a variety of factors

What often doesn't work

- Suspend to RAM
 - Two levels of suspend to RAM
 - Under S1, CPU and PSU are not turned off
 - Under S3, only power to RAM is kept
 - S1 often works even when S3 doesn't but...
 - ...S3 saves more power
 - Sometimes the machine never powers down
 - Sometimes reboots instead of suspending
 - Drivers can veto the suspend (and sometimes applications can too)

What often doesn't work

- Resuming from Suspend to RAM
 - Sometimes resumes and hangs powering up
 - Sometimes resumes and reboots
 - Sometimes resumes and doesn't turn the screen back on
 - Sometimes resumes and screen fades to white
 - Sometimes you can resume and suspend a few times before hanging
 - Certain hardware no longer works after resuming
 - Best to report this to kernel developers

What often doesn't work

- Graphics cards are problematic
 - Sometimes OS can ask the BIOS for help (three different ways of doing this)
 - Sometimes OS can save the state of graphics card to disk and restore it on resume (there are currently 10 different ways of doing this)
 - Some binary graphics drivers don't want to work with open source AGP drivers (NVIDIA...)
 - Which method works depends on graphics card and BIOS (and sometimes BIOS revision)

What often isn't configured

- Putting hardware into low power mode when “idle”
 - Turning off the wired LAN card when running from battery if there's no network cable plugged in
 - Switching wifi cards into low power mode when not in use or the signal is strong when on battery
 - Spinning down disks in laptops when on battery
 - Backlight control on non IBM laptops
 - Wake on LAN

The future

- Tickless kernel
 - Batches timer events together
 - Allows CPU to stay idle for longer saving more power
 - Needs applications to stop polling too frequently
- Fancier CPU control
- Turning off CPUs in multi CPU machines
- Better network availability handling (Network Manager)

Please help improve Linux suspend to ram on your laptop. If you find a setting that works, report the results back to the S2ram mailing list:

<http://en.opensuse.org/S2ram>

References

- The ACPI spec
 - <http://www.acpi.info/spec.htm>
- ACPI / Power Management
 - <http://www.microsoft.com/whdc/system/pnppwr/powermgmt/default.mspx>
- How Linux suspend and resume works in the ACPI age
 - <http://www.advogato.org/article/913.html>
- S2ram documentation
 - <http://en.opensuse.org/S2ram>
- Tickless kernel patches
 - <http://kerneltrap.org/node/6750>