

Power Management on DICE desktops

Chris Cooke
cc@inf.ed.ac.uk

What I'll talk about

- Power management
- What the project will deliver
- Main areas to look at
- Money

Power management

- Basically – making computers use less electricity.
- Or, better – making them use it **more efficiently**.

on PCs

- ◉ Formerly - APM
- ◉ Now - ACPI

APM

- It was fairly simple.
- It worked (eventually).

APM

- Control: APM was done at the BIOS level – the OS had little or no control.
- Hardware: the BIOS reckoned it knew exactly what hardware the machine had.
- APM worked ...
- ... but it wasn't very flexible.

ACPI

- The OS is in charge
- The ACPI spec defines multiple "states" - sleep states, CPU power states, etc.
- The OS asks the BIOS to do things

ACPI

- Very flexible
- Lots more interaction between OS and BIOS
- Lots of chances for BIOS bugs to mess things up

BIOS bugs

- According to Avery Pennarun:
- "BIOS programmers are crackheaded morons who can't implement an API correctly to save their lives."
- <http://www.advogato.org/person/apenwarr/diary.html?start=167>

ACPI

- It's far more flexible than APM.
- It's far more complicated than APM.
- It sometimes works ...
- ... sort of.

ACPI on Linux

- It's Open Source, so...
- Projects compete
- Projects fork and rename
- pm-utils, Powersave, swsusp, uswsusp, suspend, suspend2, TuxOnIce, gnome-power-manager, acpid, ...

The project

- Just a report.
- Just looking at DICE desktops.

Main Areas

Could we make our DICE desktops...

- "sleep" at night?
- operate more efficiently when running?

Sleeping at night

- What normally happens at night?
- Could it happen at other times?
- Or could we automatically wake up at night, perform maintenance (e.g. updates) then go back to sleep?

Sleeping at night

- FC6 already includes support for **suspend** and **hibernate**

(Digression)

- **suspend** – save the state to memory. Quick to suspend. Quick to resume. Needs power. Short term.
- **hibernate** – save the state to disk. No power needed – the computer can be unplugged. Takes a few seconds.

Sleeping FC6 DICE

- Sleeping for short periods - no ill effects?
- Sleeping for hours - **amd** dies on resume which freezes X and logins :-)

Efficiency

- Use less power while running.
- Get components to enter low power states for short periods (ms or even μ s).

Efficiency

- Mostly done in the kernel
- e.g. **tickless kernel** - gets rid of the Linux kernel's regular timing "tick" which does thousands of interrupts per second.

Efficiency

- Lots of possible tweaks.
- What tweaks would give us the biggest savings?
- What kernel version would we need?

Electricity costs Money



Without sleep

- My Dell GX745 uses 64W when idle.
- $64\text{W} \times 24 \text{ hours} \times 350 \text{ days} = 537.6 \text{ kWh}$
- Cost of electricity is about 10p per kWh
- So my Dell costs \sim £54 a year to power
- ... doing nothing

Without sleep

- In a Condor pool it averages 20% more than when idle = £65
- Running flat out, the machine can use about 90W (instead of 64W) = £75

While sleeping

- A suspended Dell takes 15W
- A hibernated Dell takes 12W
- A switched off Dell takes 12W
- An unplugged Dell takes 0W

Suspended overnight

- Currently - £54 to £75.
- Permanently suspended, my Dell would cost £12 a year to run.
- Suspended for 8 hours a night, £36.

Sleep savings

- We have about 600 active DICE desktops.
- Suspend them all for 8 hours a night
- $£54 \times 600 = £32,000$
- $£36 \times 600 = £22,000$
- We'd save £10,000

Condor

- Of course Condor complicates matters
- but at least we'll know roughly what it adds to our electricity bill
- which could be important for FEC.
- Future vision: Condor could wake sleeping machines as needed to clear its job queue.

Summary

- How to cope with sleeping overnight
- How to run more efficiently
- Some figures in £££