



Linux Power Management

LUG@GT 2011.09.28

Power Performance Difference between Linux Versions



Power Savings on Mobile Platform



Intel's LessWatts.org

Intel-sponsored project to improve Linux power consumption

Major projects:

- PowerTOP
- Tickless idle
- Linux ACPI
- Linux Battery Life Toolkit (BLTK)
- Power management on Intel graphics chipsets

PowerTOP

Goals of PowerTOP:

- Show how well your system is using the various hardware power-saving features
- Show you the culprit software components that are preventing optimal usage of your hardware power savings
- Help Linux developers test their application and achieve optimal behavior
- Provide you with tuning suggestions to achieve low power consumption

Tickless Idle

- Timer does a variety of things, such as process accounting, scheduler load balancing, and maintaining per-CPU timer events.
- 100Hz-1000Hz
- With "tickless idle", the Linux kernel has eliminated this periodic timer tick when the CPU is idle. This allows the CPU to remain in power saving states for a longer period of time, reducing the overall system power consumption.
- Benefits of tickless idle will be lost if the CPU is frequently awakened by unnecessary timer events.

```
zcat /proc/config.gz | grep HZ
```

CPU

CpuFreq

```
cat /sys/devices/system/cpu/cpuX/cpufreq/scaling_available_governors
```

- ondemand - Best choice for a laptop on battery
- interactive - Good choice for a laptop on AC
- conservative - Not the most power-efficient
- powersave - Minimum frequency
- performance - Like cruising down the highway in 1st

Undervolting

- PHC
- TPC

Load Distribution - for light load use 1 CPU core

```
echo 1 > /sys/devices/system/cpu/sched_mc_power_savings
```

Suspending & Hibernating

pm-utils

kernel - Built in, least flexible

uswsusp - userspace software suspend

tuxonice - 3rd party kernel module

Hard drives

- **ALPM** - Aggressive Link Power Management
 - Puts AHCI SATA link into a very low power mode when there's no IO for a while... 0.5 to 1.5 watts

```
echo "min_power" > /sys/class/scsi_host/host0/link_power_management_policy
```

- **noatime** - Don't record file access
 - Save writes to disk

```
mount -o remount,noatime /
```

- **relatime** is a good compromise

- **tmpfs** - keep temporary files in ram instead of writing to disk.

Hard drives II

- **hdparm** - Disk power management
 - -B Set power management level (1-255)
 - -S Set spindown timeout

```
hdparm -B 1 -S 12 /dev/sda #60 seconds
```

Be mindful of cycle counts!

- **VM writeback time** - Buffer disk writes
 - Lets the VM subsystem group writes for efficiency
 - Keeps disk asleep

```
echo 1500 > /proc/sys/vm/dirty_writeback_centiseecs #15 seconds
```

- **"-" in /etc/syslog.conf**
 - No sync after new log entry

Graphics

- **ASPM** - PCIe Active State Power Management
 - Might need to append `pcie_aspm=force` to kernel

```
echo "powersave" > /sys/module/pcie_aspm/parameters/policy
```

- **ATI OSS**

```
echo "dynpm" > /sys/class/drm/card0/device/power_method
```

- **ATI Catalyst**

- glakkeclock to underclock

- **NVIDIA**

- nvclock to underclock video card

Misc

Set ethernet speed to 100mbit, Wake On LAN off

```
ethtool -s eth0 autoneg off speed 100 wol d
```

Power management for Intel Wifi cards

```
iwpriv eth1 set_power 5
```

in modprobe.conf:

```
options usbcore autosuspend=1
```