

An up-to-date look at free software and its makers

# PROJECTS ON THE MOVE

A new status report on suspend and hibernate discusses the workings of, issues with, and future plans for power-saving modes in Linux.

BY CARSTEN SCHNOBER

**G**igantic pop concerts all over the world, a worried former vice president of the USA, the United Nations Climate Report – a lot of people are concerned with protecting our environment, a topic that includes saving power. Many arguments support computer users saving power, including helping the environment, saving money, or simply extending the life of laptop batteries. When it comes to power management, Linux is not generally regarded as a role model. Besides excessive power consumption during normal operations, standby functions don't always work as their makers intended.

## Saving Power

In the Linux kernel, Rafael J. Wysocki is mainly responsible for ongoing development of the Swsusp subsystem, which is

responsible for the suspend and hibernate states.

Wysocki recently composed a report to help the developer community keep abreast of the workings of these kernel components, at the same time detailing the development roadmap [1]. Wysocki wrote a similar report [2] a year ago. He gave up his initial plan of updating the report every three or four months because the Swsusp subsystem has been through so many major changes.

## ACPI

For a couple of years, the operating system-in-

dependent ACPI (Advanced Configuration and Power Interface [3]) has handled power-saving functions (and other things) on most PCs. The tasks assigned to the ACPI standard, which was developed by corporations such as Hewlett-Packard, Intel, Microsoft, Phoenix, and Toshiba, include configuring attached hardware and monitoring the system temperature and system components such as fans.

## Suspend to RAM

An assortment of non-operation states are implemented in a variety of ways by different systems.

In Suspend to RAM mode, all of the computer components, with the exception of RAM, are switched off. The CPU register content and the sta-



tus information required for restarting are written to RAM.

## Stopgrant

Although RAM obviously needs power, consumption is far below what a fully operational system would use. One variant on this mode is Stopgrant, in which the CPU remains active and keeps the data in its own registers but does not accept any instructions.

## Suspend to Disk

Suspend to Disk is a third alternative. Data is written out to disk and the computer powers off completely. When switched on, the system reads an image and restores the previous state.

## Terms

The ACPI standard defines the modes I have just described as sleep states and assigns numbers. Suspend to RAM is 4 (S4 for State 4), Suspend to Disk is 3 (S3 for State 3), and Stopgrant is S1.

Unfortunately, the terms for these states also vary. Windows calls Suspend to Disk “hibernation,” and standby mode in the world of Windows is either Suspend to RAM or Stopgrant, depending on the BIOS configuration.

On Linux, different terms are used depending on your choice of desktop environment. In kernelspeak, S1 is known as standby, S3 as suspend, and S4 is known as hibernation.

## Stumbling Blocks

ACPI was designed to make everything simple by offloading much of the work from the operating system. In an ideal world, systems with ACPI would receive the same instructions and then switch the individual components to the state desired by the user. Which operating system sent the instructions would not make any difference.

In the real world, various issues prevent power management from being that simple. One problem is that some mainboards do not implement ACPI correctly at BIOS level. Many manufacturers are obviously content if things work on Windows. In some cases, proprietary drivers translate instructions to allow the board to trigger sleep states despite the lack of an ACPI implementation.

Of course, drivers are not available for Linux, and the idea of fixing a broken

implementation by writing more software is unlikely to impress open source developers and users. This issue prompted kernel programmers to develop their own approach to putting computers to sleep – software suspend.

## Software Suspend

Software suspend is largely independent of both ACPI and the motherboard and leaves the task of restoring the previous state when the computer wakes up to the corresponding kernel modules; however, particularly in the case of Suspend to RAM, the issue affects more than just the mainboard. The system has to ascertain the status of a whole bunch of components, including the CPU, the graphics card, USB, and other hardware, along with its own status and that of the device drivers and the BIOS, to store an image of the current system state in RAM for restoration later. For an error-free restore, device drivers all need to support error-free translation of sleep and wake-up events for the hardware in question.

Things aren’t looking too good if you take a look at the current performance of power-saving modes on Linux, although much progress has been made in the past few years. Some hardware drivers still fail to cooperate, and whether power saving works will depend on the hardware in question.

Laptop manufacturers tend to keep to the standards, which is good because this is where you really need power-saving functionality, and this means that you stand a good chance that the various suspend modes will work on your laptop. Most systems support Suspend to Disk anyway because it demands less of components and their respective drivers.

## Future

Wysocki sees device and driver handling as the main problem for the Swsusp subsystem because it is not the operating system’s task to resolve issues caused by

```

jcasad@sugar: ~
File Edit View Terminal Tabs Help
[ 1211.632000] ohci1394 0000:02:09.1: LATE suspend
[ 1211.632000] yenta_cardbus 0000:02:09.0: LATE suspend
[ 1211.632000] pci 0000:02:07.0: LATE suspend
[ 1211.632000] pci 0000:02:05.0: LATE suspend
[ 1211.632000] pci 0000:00:1f.3: LATE suspend
[ 1211.632000] pci 0000:00:1f.0: LATE suspend
[ 1211.632000] Intel ICH Modem 0000:00:1e.3: LATE suspend
[ 1211.632000] Intel ICH 0000:00:1e.2: LATE suspend
[ 1211.632000] pci 0000:00:1e.0: LATE suspend
[ 1211.632000] pciport-driver 0000:00:1c.0: LATE suspend
[ 1211.632000] pci 0000:00:02.1: LATE suspend
[ 1211.632000] pci 0000:00:02.0: LATE suspend
[ 1211.632000] agpgart-intel 0000:00:00.0: LATE suspend
[ 1211.632000] Back to CI
[ 1218.632000] agpgart-intel 0000:00:00.0: EARLY resume
[ 1218.632000] pci 0000:00:02.0: EARLY resume
[ 1218.632000] pci 0000:00:02.1: EARLY resume
[ 1218.632000] pciport-driver 0000:00:1c.0: EARLY resume
[ 1218.632000] uhci_hcd 0000:00:1d.0: EARLY resume
[ 1218.632000] uhci_hcd 0000:00:1d.1: EARLY resume
[ 1218.632000] uhci_hcd 0000:00:1d.2: EARLY resume
[ 1218.632000] uhci_hcd 0000:00:1d.3: EARLY resume
[ 1218.632000] ehci_hcd 0000:00:1d.7: EARLY resume
[ 1218.632000] pci 0000:00:1e.0: EARLY resume
[ 1218.632000] Intel ICH 0000:00:1e.2: EARLY resume

```

Figure 1: Before going to suspend, the system switches off every device, and during the resume it switches them on again.

faulty ACPI implementations. The next step toward improving the status quo is disambiguating hibernation and suspend, which have used the same system calls thus far.

Various issues with freezing and continuing certain tasks also still exist. The hibernate or suspend procedure is canceled if it proves impossible to freeze a task. The alternative would be to switch the system to a non-active state despite the problem; however, the risk of the hardware not working when the system wakes up again is far too high. For Wysocki, bailing out at the last minute is the lesser of two evils. Of course, this isn’t a genuine solution, so Wysocki will be looking to find a different one in the near future.

## TuxOnIce

A second project, dubbed TuxOnIce [4] and led by Nigel Cunningham, focuses on the same task but is not part of the official Linux kernel. The package, previously called Suspend2, includes additional features, such as the ability to store system state images in files rather than in swap partitions and the ability to encrypt images. ■

## INFO

- [1] Status report on suspend and hibernation 2007:  
<http://article.gmane.org/gmane.linux.power-management.general/7701>
- [2] Swsusp status report 2006:  
<http://lkml.org/lkml/2006/7/25/105>
- [3] ACPI: <http://www.acpi.info>
- [4] TuxOnIce: <http://www.tuxonice.net>