

Re: Flames over -- Re: Which is simpler?

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-02/msg04563.html>

- *From:* Phillip Susi <psusi@xxxxxxxxxx>
 - *Date:* Mon, 13 Feb 2006 12:14:55 -0500
-

Alan Stern wrote:

Sorry, but you're wrong. First of all, testing if hardware is there is different from testing whether it has changed -- it could have changed while the system was asleep, with the result that hardware is indeed there but it's not the `_same_` hardware.

If you believe I am wrong, then please offer proof. Specifically, refer to the kernel code that handles verifying that the hardware is still there after a resume. I very well may be wrong, but you can not simply surmise this to be so without any proof. It is currently my belief that the kernel probes the hardware the same way after a cold boot, resume from `s-t-r`, and resume from `s-t-d`. Specifically it looks to see if a device is located in the same bus location with the same serial number, etc, and if so, access to it continues exactly where it left off prior to the suspend.

This is why you can suspend to disk, and when you resume, your open files are still valid; the disk is found to be still there, so it continues to be accessible. Nothing gets clobbered as a result of the disk seeming to be disconnected and reconnected. If that does happen, it is a bug.

Second, with USB at any rate, in addition to checking that hardware is still there, the kernel queries the USB controller to see if a disconnect occurred while the system was asleep. (If the controller wasn't powered during that time then it will report that every USB device was disconnected.)

AFAIK, there is no interface by which the kernel can query that information from the controller, maybe you could show me? If that is the case however, then I consider that to be a bug with the USB bus and the kernel's handling of it. The kernel needs to be able to assume that nothing was disconnected while it was shutdown, provided that the same devices are there now as when it went to sleep. This is how it behaves for say, SCSI disks in desktops/servers, where the controller certainly is completely powered off. It should work the same for USB.

Third, there is indeed something special that monitors USB device insertion/removal while suspended to RAM -- the USB host controller does so if it has suspend power.

Could you site references to that? AFAIK, the host controller is only capable of generating a wake event when

Re: Flames over -- Re: Which is simpler?

the bus state changes; it does not have a means of informing the OS what has changed, aside from the OS enumerating the devices on the bus again.

No. It does work exactly as designed and it's not buggy. You just don't understand it.

If the mounted filesystem becomes corrupted over a hibernation because the kernel thinks the drive was unplugged, then plugged back in, that clearly is a bug. It does not do this for disks connected via other bus types, and it clearly is undesirable to corrupt data in this way.

The state of the kernel after resuming from either suspend to disk, or suspend to ram is the same. The filesystem is still mounted, and any dirty pages in ram will be flushed just like normal. Whether the disk is connected via SCSI, SATA, USB, or whatever does not matter.

Don't be silly. Dirty pages can't be flushed to disks that are no longer attached! And if a USB disk was unplugged while the system was asleep, the kernel will know that it is no longer attached. I don't know which other bus drivers check for this sort of thing.

The disk is still attached of course. The scenario we are talking about does not actually involve disconnecting the drive. Since the drive is not actually disconnected, if the kernel believes it has been, it's a bug.

The kernel knows that the same device is still there on the bus, and so it picks up right where it left off. If it thinks the device has been unplugged and reconnected, that is a bug.

It's not a bug if the device `_has_` been unplugged and reconnected. When that happens, there's no way for the kernel to tell whether the device there now is the same as the device that used to be there.

Again, we're not talking about it actually being unplugged, though since the kernel has no way of knowing it, you can unplug a scsi disk while hibernated, then plug it back in before you resume, and it will work just fine since the same type of hardware with the same serial number et al is found in the same place on the bus.

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>