

[patch 1/1] cpufreq_ondemand/conservative: invert meaning of 'ignore_nice'

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-10/1709.html>

From: Alexander Clouter (alex-kernel_at_digriz.org.uk)

Date: 10/07/05

Date: Fri, 7 Oct 2005 09:09:44 +0100

To: LKML <linux-kernel@vger.kernel.org>, cpufreq@lists.linux.org.uk

The use of the 'ignore_nice' sysfs file is confusing to anyone using. This patch makes it so when you now set it to the default value of 1, process nice time is also ignored in the cpu 'busyness' calculation. We also rename 'ignore_nice' to 'ignore_nice_task' so that userland tools notice the inverse behaviour sooner.

Prior to this patch to set it to '1' to make process nice time count...even confused me :)

WARNING: this obvious breaks any userland tools that expect things to be the other way round; hence why, as suggested, the rename of the 'ignore_nice' file to 'ignore_nice_tasks'

Signed-off-by: Alexander Clouter <alex-kernel@digriz.org.uk>

—

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

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

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

• text/plain attachment: [01_inverse_ignore_nice_flag.diff](#)

• application/pgp-signature attachment: [Digital signature](#)