

## Queues when accessing disks

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-12/6355.html>

---

**From:** Felipe Erias ([charles.swann\\_at\\_gmail.com](mailto:charles.swann_at_gmail.com))

**Date:** 12/31/04

Date: Fri, 31 Dec 2004 04:32:59 +0100

To: [linux-kernel@vger.kernel.org](mailto:linux-kernel@vger.kernel.org)

Hi,

I'm trying to apply queuing theory to the study of the GNU/Linux kernel. Right now, I'm focusing in the queue of processes that appears when they try to access an I/O device (specifically, an IDE HD). When they want to read data, it behaves as a usual queue: several clients (processes) that require attention from a server (disk / driver / ...). The case when they want to write data is a bit more tricky, because of the cache buffers used by the OS, and maybe could be modeled by a network of queues. Both cases are interesting for my work, but I'll take the reading one first, just because it seems a bit more simple 'a priori'.

To modelize the queue, I need to get some information:

- what processes claim attention from the disk
- when they do it
- when they begin to be served
- when they finish being served

To get all this information, maybe I could hack my kernel a bit to write a line to a log on every access to the HD, or account the IRQs from the IDE channels... I also have the feeling that this queuing problem could dissappear o became more hidden if DMA were enabled.

To be true, I'm a bit lost and that's why I ask for your help.

Yours sincerely,

Felipe Erias

-

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