Re: A scrub daemon (prezeroing)

Source: http://linux.derkeiler.com/Mailing−Lists/Kernel/2005−02/0575.html

---

From: Maciej W. Rozycki (*macro_at_linux−mips.org*)
Date: 02/02/05

```
Date:   Wed, 2 Feb 2005 21:00:46 +0000 (GMT)
To: Marcelo Tosatti <marcelo.tosatti@cyclades.com>
```

On Wed, 2 Feb 2005, Marcelo Tosatti wrote:

> > Some architectures tend to have spare DMA engines lying around. There's
> > no need to use the CPU for zeroing pages. How feasible would it be for
> > scrubd to use these?
[...]
> I suppose you are talking about DMA engines which are not being driven
> by any driver ?

E.g. the Broadcom's MIPS64−based SOCs have four general purpose DMA engines onchip which can transfer data to/from the memory controller in 32−byte chunks over the 256−bit internal bus. We have hardly any use for these devices and certainly not for all four of them.

> Sounds very interesting idea to me. Guess it depends on whether the cost of
> DMA write for memory zeroing, which is memory architecture/DMA engine dependant,
> offsets the cost of CPU zeroing.

I suppose so, at least with the Broadcom's chips you avoid cache trashing, yet you don't need to care about stale data as coherency between CPUs and the onchip memory controller is maintained automatically by hardware.

Maciej
−