

---

Subject: Re: [PATCH v3 3/6] expose fine-grained per-cpu data for cpuacct stats  
Posted by [Glauber Costa](#) on Wed, 30 May 2012 12:52:31 GMT

[View Forum Message](#) <> [Reply to Message](#)

---

On 05/30/2012 04:48 PM, Paul Turner wrote:

> a) the information in /proc/stat is actually much denser since it's  
> "cpu VAL VAL VAL VAL" as opposed to "cpuX.FIELD VAL"

easily fixable here. Less descriptive, but we can use a header line with the description much like how /proc/slabinfo does, and we still have an extensible interface that is dense, at the same time.

> b) If it became a problem the /proc/stat case is actually fairly  
> trivially fixable by defining each cpu as a record and "everything  
> else" as a magic in-out-of-cpus value.

>

>> >

>> > Now, if you guys are okay with a file per-cpu, I can do it as well.

>> > It pollutes the filesystem, but at least protects against the fact that this

>> > is kmalloc-backed.

>> >

> As I prefaced, I'm not sure there's much that can be trivially done

> about it. This is really a fundamental limitation of how read\_map()

> works.

>

> What we really need is a proper seq\_file exposed through cftypes.

That can be done.

---