
Subject: Re: [ckrm-tech] [PATCH] BC: resource beancounters (v4) (added user memory)

Posted by [Dave Hansen](#) on Fri, 08 Sep 2006 15:30:50 GMT

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

On Tue, 2006-09-05 at 17:17 -0700, Rohit Seth wrote:

> I'm wondering why not have different processes to serve different
> domains on the same physical server...particularly when they have
> different database to work on.

This is largely because this is I think how it is done today, and it has a lot of disadvantages. They also want to be able to account for traffic on the same database. Think of a large web hosting environment where you charged everyone (hundreds or thousands of users) by CPU and I/O bandwidth used at all levels of a given transaction.

> Is the amount of memory that you save by
> having a single copy that much useful that you are even okay to
> serialize the whole operation (What would happen, while the request for
> foo.com is getting worked on, there is another request for
> foo_bar.com...does it need to wait for foo.com request to get done
> before it can be served).

Let's put it this way. Enterprise databases can be memory pigs. It isn't feasible to run hundreds or thousands of copies on each machine.

-- Dave
