
Subject: Re: [RFC][PATCH] Per container statistics
Posted by [Balbir Singh](#) on Tue, 15 May 2007 07:23:00 GMT
[View Forum Message](#) <> [Reply to Message](#)

Kirill Korotaev wrote:

> oh, please no... Andrew, this loop can be very very long when having > 10,000 tasks on the machine...
> we have had enough such issues in OpenVZ and just don't want to come through this again.
> Also sum of RUNNING + UNINTERRUPTIBLE is required for per-container loadavg calculations.
>

Hi, Kirill,

I've asked Paul Menage to help us with this scenario by adding a per-container task list. Would you be ok, if we walked the per-container task list. We could further optimize the routine by adding

1. A container id (avoid parsing file paths to uniquely identify a container)
2. If we cannot do (1), then cache container paths

The other option is of-course what Pavel posted, but I think the amortized cost of tracking task state at the time of task state change or looping over tasks on demand should be the same.

Do you agree in general with containerstats? How we build (walk the task list or at the time of task state change) is flexible with containerstats.

Thanks for your comments!

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL
