
Subject: Re: [ckrm-tech] [RFC] [PATCH 00/12] CKRM after a major overhaul
Posted by [Kirill Korotaev](#) on Mon, 24 Apr 2006 06:38:30 GMT

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

>>>>Yes, it is effective, and the reclamation is O(1) too. It has couple of
>>>>problems by design, (1) doesn't handle shared pages and (2) doesn't
>>>>provide support for both min_shares and max_shares.
>>>
>>>Right. I wanted to show proof-of-concept of the pzone based controller
>>>and implemented minimal features necessary as the memory controller.
>>>So, the pzone based controller still needs development and some cleanup.
>>
>>Just out of curiosity, how it was measured that it is effective?
>
>
> I don't have any benchmark numbers yet, so I can't explain the
> effectiveness with numbers. I've been looking for the way to
> measure the cost of pzones correctly, but I've not found it out yet.
>
>
>>How does it work when there is a global memory shortage in the system?
>
>
> I guess you are referring to the situation that global memory is running
> out but there are free pages in pzones. These free pages in pzones are
> handled as reserved for pzone users and not used even in global memory
> shortage.
ok. Let me explain what I mean.
Imagine the situation with global memory shortage. In kernel, there are
threads which do some job behalf the user, e.g. kjournald, loop etc. If
the user has some pzone memory, but these threads fail to do their job
some nasty things can happen (ext3 problems, deadlocks, OOM etc.)
If such behaviour is ok for you, then great. But did you consider it?

Also, I can't understand how it works with OOM killer. If pzones has
enough memory, but there is a global shortage, who will be killed?

Thanks,
Kirill

Subject: Re: [ckrm-tech] [RFC] [PATCH 00/12] CKRM after a major overhaul
Posted by [KUROSAWA Takahiro](#) on Mon, 24 Apr 2006 07:12:46 GMT

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

On Mon, 24 Apr 2006 10:45:59 +0400
Kirill Korotaev <dev@openvz.org> wrote:

> >>>>Yes, it is effective, and the reclamation is O(1) too. It has couple of
> >>>>problems by design, (1) doesn't handle shared pages and (2) doesn't
> >>>>provide support for both min_shares and max_shares.
> >>>>
> >>>>Right. I wanted to show proof-of-concept of the pzone based controller
> >>>>and implemented minimal features necessary as the memory controller.
> >>>>So, the pzone based controller still needs development and some cleanup.
> >>>>
> >>>>Just out of curiosity, how it was measured that it is effective?
> >>>>
> >>>>I don't have any benchmark numbers yet, so I can't explain the
> >>>>effectiveness with numbers. I've been looking for the way to
> >>>>measure the cost of pzones correctly, but I've not found it out yet.
> >>>>
> >>>>How does it work when there is a global memory shortage in the system?
> >>>>
> >>>>I guess you are referring to the situation that global memory is running
> >>>>out but there are free pages in pzones. These free pages in pzones are
> >>>>handled as reserved for pzone users and not used even in global memory
> >>>>shortage.
> >>>>ok. Let me explain what I mean.
> >>>>Imagine the situation with global memory shortage. In kernel, there are
> >>>>threads which do some job behalf the user, e.g. kjournald, loop etc. If
> >>>>the user has some pzone memory, but these threads fail to do their job
> >>>>some nasty things can happen (ext3 problems, deadlocks, OOM etc.)
> >>>>If such behaviour is ok for you, then great. But did you consider it?
> >>>>
> >>>>Also, I can't understand how it works with OOM killer. If pzones has
> >>>>enough memory, but there is a global shortage, who will be killed?

I understand.

IMHO, only the system processes should use global memory.
User processes that may cause such memory shortage should be
enclosed in pzones first.

Thanks,

--

KUROSAWA, Takahiro
