
Subject: Re: [PATCH 0/4] swapcgroup(v2)
Posted by [Balbir Singh](#) on Fri, 23 May 2008 06:45:15 GMT
[View Forum Message](#) <> [Reply to Message](#)

KOSAKI Motohiro wrote:

>>> Have you seen any real world example of this?
>> At the unsophisticated end, there are lots of (Fortran) HPC applications
>> with very large static array declarations but only "use" a small fraction
>> of that. Those users know they only need a small fraction and are happy
>> to volunteer small physical memory limits that we (admins/queuing
>> systems) can apply.
>>
>> At the sophisticated end, the use of numerous large memory maps in
>> parallel HPC applications to gain visibility into other processes is
>> growing. We have processes with VSZ > 400GB just because they have
>> 4GB maps into 127 other processes. Their physical page use is of
>> the order 2GB.
>
> Ah, agreed.
> Fujitsu HPC user said similar things ago.

OK, so this use case is HPC specific. I am not against the swap controller, but overcommit can lead to problems if not controlled - such as OOM kill. The virtual address space limit helps applications fail gracefully rather than swap out excessively or OOM.

I suspect there'll be applications that swing both ways.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL

Containers mailing list
Containers@lists.linux-foundation.org
<https://lists.linux-foundation.org/mailman/listinfo/containers>
