
Subject: Re: [PATCH 0/4] swapcgroup(v2)

Posted by [David Singleton](#) on Fri, 23 May 2008 05:29:31 GMT

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

Balbir Singh wrote:

> KOSAKI Motohiro wrote:

>>> One option is to limit the virtual address space usage of the cgroup to ensure
>>> that swap usage of a cgroup will *not* exceed the specified limit. Along with a
>>> good swap controller, it should provide good control over the cgroup's memory usage.

>> unfortunately, it doesn't work in real world.

>> IMHO you said as old good age.

>>

>> because, Some JavaVM consume crazy large virtual address space.

>> it often consume >10x than physical memory consumption.

>>

>

> 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.

Imposing virtual address space limits on these applications is meaningless.

The overcommit feature of Linux.

> We usually by default limit the overcommit to 1.5 times total memory (IIRC).

> Yes, one can override that value, you get the same flexibility with the virtual
> address space controller.

>

> I thought java was particular about it with its heap management options and policy.

>

>> yes, that behaviour is crazy. but it is used widely.

>> thus, We shouldn't assume virtual address space limitation.

>

> It's useful in many cases to limit the virtual address space - to allow
> applications to deal with memory failure, rather than

>

> 1. OOM the application later

> 2. Allow uncontrolled swapping (swap controller would help here)

>

David

Containers mailing list

Containers@lists.linux-foundation.org

<https://lists.linux-foundation.org/mailman/listinfo/containers>
