Subject: Re: [PATCH v5 14/14] Add documentation about the kmem controller
Posted by Michal Hocko on Tue, 16 Oct 2012 12:23:24 GMT
View Forum Message <> Reply to Message

On Tue 16-10-12 14:16:51, Glauber Costa wrote:
> Signed-off-by: Glauber Costa <glommer@parallels.com>
> CC: Frederic Weisbecker <fweisbec@redhat.com>
> CC: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
> CC: Michal Hocko <mhocko@suse.cz>
> CC: Christoph Lameter <cl@linux.com>
> CC: Pekka Enberg <penberg@cs.helsinki.fi>
> CC: Johannes Weiner <hannes@cmpxchg.org>
> CC: Suleiman Souhlal <suleiman@google.com>
> CC: Tejun Heo <tj@kernel.org>

Acked-by: Michal Hocko <mhocko@suse.cz

Just a nit..
> ---
> Documentation/cgroups/memory.txt | 58 ++++++++++++++++++++++++++++++++++++++-
> 1 file changed, 57 insertions(+), 1 deletion(-)
>
> diff --git a/Documentation/cgroups/memory.txt b/Documentation/cgroups/memory.txt
> index c07f7b4..dd15be8 100644
> --- a/Documentation/cgroups/memory.txt
> +++ b/Documentation/cgroups/memory.txt
[...]
> @@ -268,20 +273,65 @@ the amount of kernel memory used by the system. Kernel memory is fundamentally
> different than user memory, since it can't be swapped out, which makes it
> possible to DoS the system by consuming too much of this precious resource.
>
> +Kernel memory won't be accounted at all until limit on a group is set. This
> +allows for existing setups to continue working without disruption. The limit
> +cannot be set if the cgroup have children, or if there are already tasks in the
> +cgroup. When use_hierarchy == 1 and a group is accounted, its children will
> +automatically be accounted regardless of their limit value.
> +
> +After a controller is first limited, it will be kept being accounted until it

s/controller/group/

> +is removed. The memory limitation itself, can of course be removed by writing
> +-1 to memory.kmem.limit_in_bytes. In this case, kmem will be accounted, but not
> +limited.
> +

--

Michal Hocko
SUSE Labs