

On Tue, 19 Sep 2006, Rohit Seth wrote:

> +Currently we are tracking user memory (both file based
> +and anonymous). The memory handler is currently deactivating pages
> +belonging to a container that has gone over the limit. Even though this
> +allows containers to go over board their limits but 1- once they are
> +over the limit then they run in degraded manner and 2- if there is any
> +memory pressure then the (extra) pages belonging to this container are
> +the prime candidates for swapping (for example). The statistics that
> +are shown in each container directory are the current values of each
> +resource consumption.

Containers via cpusets allow a clean implementation of a restricted memory area. The system will swap and generate an OOM message if no memory can be recovered.

> +4- Add a task to container
> + cd /mnt/configfs/cotainers/test_container
> + echo <pid> > addtask
> +
> +Now the <pid> and its subsequently forked children will belong to container
> +test_container.
> +
> +5- Remove a task from container
> + echo <pid> > rmtask

Could you make that compatible with the way cpusets do it?

> +9- Freeing a container
> + cd /mnt/configfs/containers/
> + rmdir test_container

Adding and removal is the same way as cpusets.
