

Hi,

What I am proposing with this series is a stripped down version of the kmem controller for memcg that would allow us to merge significant parts of the infrastructure, while leaving out, for now, the polemic bits about the slab while it is being reworked by Cristoph.

Me reasoning for that is that after the last change to introduce a gfp flag to mark kernel allocations, it became clear to me that tracking other resources like the stack would then follow extremely naturally. I figured that at some point we'd have to solve the issue pointed by David, and avoid testing the Slab flag in the page allocator, since it would soon be made more generic. I do that by having the callers to explicit mark it.

So to demonstrate how it would work, I am introducing a stack tracker here, that is already a functionality per-se: it successfully stops fork bombs to happen. (Sorry for doing all your work, Frederic =p). Note that after all memcg infrastructure is deployed, it becomes very easy to track anything. The last patch of this series is extremely simple.

The infrastructure is exactly the same we had in memcg, but stripped down of the slab parts. And because what we have after those patches is a feature per-se, I think it could be considered for merging.

Let me know what you think.

Glauber Costa (9):

- memcg: change defines to an enum
- kmem slab accounting basic infrastructure
- Add a __GFP_KMEMCG flag
- memcg: kmem controller infrastructure
- mm: Allocate kernel pages to the right memcg
- memcg: disable kmem code when not in use.
- memcg: propagate kmem limiting information to children
- memcg: allow a memcg with kmem charges to be destructed.
- protect architectures where THREAD_SIZE >= PAGE_SIZE against fork bombs

Suleiman Souhlal (2):

- memcg: Make it possible to use the stock for more than one page.
- memcg: Reclaim when more than one page needed.

```
include/linux/gfp.h      | 11 +-  
include/linux/memcontrol.h | 46 +++++
```

```
include/linux/thread_info.h | 6 +
kernel/fork.c                | 4 +-
mm/memcontrol.c              | 395 ++++++
mm/page_alloc.c              | 27 +++
6 files changed, 464 insertions(+), 25 deletions(-)
```

--
1.7.10.2
