
Subject: [PATCH v2 00/13] Memcg Kernel Memory Tracking.
Posted by [Suleiman Souhlal](#) on Fri, 09 Mar 2012 20:39:03 GMT
[View Forum Message](#) <> [Reply to Message](#)

This is v2 of my kernel memory tracking patchset for memcg.

Lots of changes based on feedback from Glauber and Kamezawa.
In particular, I changed it to be opt-in instead of opt-out:
In order for a slab type to be tracked, it has to be marked with
SLAB_MEMCG_ACCT at kmem_cache_create() time.
Currently, only dentries and kmalloc are tracked.

Planned for v3:

- Slub support.
- Using a static_branch to remove overhead when no cgroups have been created.
- Getting rid of kmem_cache_get_ref/drop_ref pair in kmem_cache_free.

Detailed change list from v1 (<http://marc.info/?l=linux-mm&m=133038361014525>):

- Fixed misspelling in documentation.
- Added flags field to struct mem_cgroup.
- Moved independent_kmem_limit into flags.
- Renamed kmem_bytes to kmem.
- Divided consume_stock changes into two changes.
- Fixed crash at boot when not every commit is applied.
- Moved the new fields in kmem_cache into their own struct.
- Got rid of SLAB_MEMCG slab flag.
- Dropped accounting to root.
- Added css_id into memcg slab name.
- Changed memcg cache creation to always be deferred to workqueue.
- Replaced bypass_bytes with overcharging the cgroup.
- Got rid of #ifdef CONFIG_SLAB from memcontrol.c.
- Got rid of __GFP_NOACCOUNT, changing to an opt-in model.
- Remove kmem limit when turning off independent limit.
- Moved the accounting of kmalloc to its own patch.
- Removed useless parameters from memcg_create_kmem_cache().
- Get a ref to the css when enqueueing cache for creation.
- increased MAX_KMEM_CACHE_TYPES to 400.

Suleiman Souhlal (13):

- memcg: Consolidate various flags into a single flags field.
- memcg: Kernel memory accounting infrastructure.
- memcg: Uncharge all kmem when deleting a cgroup.
- memcg: Make it possible to use the stock for more than one page.
- memcg: Reclaim when more than one page needed.
- slab: Add kmem_cache_gfp_flags() helper function.
- memcg: Slab accounting.
- memcg: Make dentry slab memory accounted in kernel memory accounting.

memcg: Account for kmalloc in kernel memory accounting.
memcg: Track all the memcg children of a kmem_cache.
memcg: Handle bypassed kernel memory charges.
memcg: Per-memcg memory.kmem.slabinfo file.
memcg: Document kernel memory accounting.

Documentation/cgroups/memory.txt | 44 +++-
fs/dcache.c | 4 +-
include/linux/memcontrol.h | 30 ++-
include/linux/slab.h | 56 ++++
include/linux/slab_def.h | 79 +++++-
include/linux/slob_def.h | 6 +
include/linux/slub_def.h | 9 +
init/Kconfig | 2 +-
mm/memcontrol.c | 633 ++++++-----
mm/slab.c | 431 ++++++-----
10 files changed, 1183 insertions(+), 111 deletions(-)

-- Suleiman
