Subject: [PATCH v5 18/18] Add slab-specific documentation about the kmem controller
Posted by Glauber Costa on Fri, 19 Oct 2012 14:20:42 GMT

View Forum Message <> Reply to Message

Signed-off-by: Glauber Costa <glommer@parallels.com>
CC: Christoph Lameter <cl@linux.com>
CC: Pekka Enberg <penberg@cs.helsinki.fi>
CC: Michal Hocko <mhocko@suse.cz>
CC: Kamezawa Hiroyuki <kamezawa.hiroyu@jp.fujitsu.com>
CC: Johannes Weiner <hannes@cmpxchg.org>
CC: Suleiman Souhlal <suleiman@google.com>
CC: Tejun Heo <tj@kernel.org>

 Documentation/cgroups/memory.txt | 7 +++++++
 1 file changed, 7 insertions(+)

diff --git a/Documentation/cgroups/memory.txt b/Documentation/cgroups/memory.txt
index 206853b..9d9938d 100644
--- a/Documentation/cgroups/memory.txt
+++ b/Documentation/cgroups/memory.txt
@@ -301,6 +301,13 @@ to trigger slab reclaim when those limits are reached.
 kernel memory, we prevent new processes from being created when the kernel
 memory usage is too high.

+* slab pages: pages allocated by the SLAB or SLUB allocator are tracked. A copy
+of each kmem_cache is created everytime the cache is touched by the first time
+from inside the memcg. The creation is done lazily, so some objects can still be
+skipped while the cache is being created. All objects in a slab page should
+belong to the same memcg. This only fails to hold when a task is migrated to a
+different memcg during the page allocation by the cache.
+
 * sockets memory pressure: some sockets protocols have memory pressure
 thresholds. The Memory Controller allows them to be controlled individually
 per cgroup, instead of globally.
--
1.7.11.7