Subject: [PATCH][for -mm] per-zone and reclaim enhancements for memory controller take 3 [0/10] introduction
Posted by KAMEZAWA Hiroyuki on Tue, 27 Nov 2007 02:54:41 GMT
View Forum Message <> Reply to Message

Hi, this is per-zone/reclaim support patch set for memory controller (cgroup).

Major changes from previous one is
 -- tested with 2.6.24-rc3-mm1 + ia64/NUMA
 -- applied comments.

I did small test on real NUMA machine.
My machine was ia64/8CPU/2Node NUMA. I tried to complile the kernel under 800M bytes limit with 32 parallel make. (make -j 32)

 - 2.6.24-rc3-mm1 (+ scsi fix) .... shows soft lock-up.
   before soft lock-up, %sys was almost 100% in several times.

 - 2.6.24-rc3-mm1 (+ scsi fix) + this set .... completed succesfully
   It seems %iowait dominates the total performance.
   (current memory controller has no background reclaim)

Seems this set give us some progress.

(*) I'd like to merge YAMAMOTO-san's background page reclaim for memory
    controller before discussing about the number of performance.

Andrew, could you pick these up to -mm ?

Patch series brief description:

[1/10] ... add scan_global_lru() macro (clean up)
[2/10] ... nid/zid helper function for cgroup
[3/10] ... introduce per-zone object for memory controller and add
          active/inactive counter.
[4/10] ... calculate mapper_ratio per cgroup (for memory reclaim)
[5/10] ... calculate active/inactive imbalance per cgroup (based on [3])
[6/10] ... remember reclaim priority in memory controller
[7/10] ... calculate the number of pages to be reclaimed per cgroup

[8/10] ... modifies vmscan.c to isolate global-lru-reclaim and
          memory-cgroup-reclaim in obvious manner.
          (this patch uses functions defined in [4 - 7])
[9/10] ... implement per-zone-lru for cgroup (based on [3])
[10/10] ... implement per-zone lru lock for cgroup (based on [3][9])

Any comments are welcome.

Thanks,
-Kame

_____

Containers mailing list
Containers@lists.linux-foundation.org
https://lists.linux-foundation.org/mailman/listinfo/containers