

On Wed, 2006-09-27 at 12:50 -0700, Chandra Seetharaman wrote:

- > Rohit,
- >
- > I finally looked into your memory controller patches. Here are some of
- > the issues I see:
- > (All points below are in the context of page limit of containers being
- > hit and the new code starts freeing up pages)
- >
- > 1. LRU is ignored totally thereby thrashing the working set (as pointed
- > by Peter Zijlstra).

As the container goes over the limit, this algorithm deactivates some of the pages. I agree that the logic to find out the correct pages to deactivate needs to be improved. But the idea is that these pages go in front of inactive list so that if there is any memory pressure system wide then these pages can easily be reclaimed.

- > 2. Frees up file pages first when hitting the page limit thereby making
- > vm_swappiness ineffective.

Not sure if I understood this part correctly. But the choice when the container goes over its limit is between swap out some of the anonymous memory first or writeback some of the dirty file pages belonging to this container.

- > 3. Starts writing back pages when the # of file pages is close to the
- > limit, thereby breaking the current writeback algorithm/logic.

That is done so as to ensure processes belonging to container (Whose limit is hit) are the first ones getting penalized. For example, if you run a tar in a container with 100MB limit then the dirty file pages will be written back to disk when 100MB limit is hit). Though I will be adding a HARD_LIMIT on page cache flag and the strict limit will be only maintained if this container flag is set.

- > 4. MAPPED files are not counted against the page limit. why ?. This
- > affects reclamation behavior and makes vm_swappiness ineffective.

num_mapped_pages only indicates how many page cache pages are mapped in user page tables. More of an accounting variable.

- > 5. Starts freeing up pages from the first task or the first file in the
- > linked list. This logic unfairly penalizes the early members of the
- > list.

This is the part that I've to fix. Some per container variables that remembers the last values will help here.

- > 6. Both active and inactive pages use physical pages. But, the
- > controller only counts active pages and not inactive pages. why ?

The thought is, it is okay for containers to go over its limit as long as there is enough memory in the system. When there is any memory pressure then the inactive (+ dereferenced) pages get swapped out thus penalizing the container. I'm also thinking of having hard limit for anonymous pages beyond which the container will not be able to grow its anonymous pages.

- > 7. Page limit is checked against the sum of (anon and file pages) in
- > some places and against active pages at some other places. IMO, it
- > should be always compared to the same value.
- >

It is checked against sum of anon+file pages at the time when new pages is getting allocated. But as the reclaimer activate the pages, so it is also important to make sure the number of active pages is not going above its limit.

Thanks for your comments,
-rohit
