

---

Subject: Re: [PATCH 8/8] Per-container pages reclamation  
Posted by [Balbir Singh](#) on Wed, 02 May 2007 09:51:03 GMT  
[View Forum Message](#) <> [Reply to Message](#)

---

Pavel Emelianov wrote:

```
> Implement try_to_free_pages_in_container() to free the
> pages in container that has run out of memory.
>
> The scan_control->isolate_pages() function isolates the
> container pages only.
>
> +#ifdef CONFIG_RSS_CONTAINER
> +unsigned long try_to_free_pages_in_container(struct rss_container *cnt)
> +{
> + struct scan_control sc = {
> + .gfp_mask = GFP_KERNEL,
> + .may_writepage = 1,
> + .swap_cluster_max = 1,
> + .may_swap = 1,
> + .swappiness = vm_swappiness,
> + .cnt = cnt,
> + .isolate_pages = isolate_pages_in_container,
> + };
> + int node;
> + struct zone **zones;
> +
> + for_each_node(node) {
> +#ifdef CONFIG_HIGHMEM
> + zones = NODE_DATA(node)->node_zonelist[ZONE_HIGHMEM].zones;
> + if (do_try_to_free_pages(zones, sc.gfp_mask, &sc))
> + return 1;
> +#endif
> + zones = NODE_DATA(node)->node_zonelist[ZONE_NORMAL].zones;
> + if (do_try_to_free_pages(zones, sc.gfp_mask, &sc))
> + return 1;
> + }
> + return 0;
> +}
> +#endif
> +
```

Hi, Pavel,

This patch fixes a bug in the RSS controller, where we walk through all nodes during reclaim in `try_to_free_pages_in_container()`. Instead of `for_each_node()`, we now use `for_each_online_node()` so that we do not end up with invalid zones from nodes that are not online.

Signed-off-by: Balbir Singh <balbir@linux.vnet.ibm.com>  
Signed-off-by: Vaidyanathan Srinivasan <svaidy@linux.vnet.ibm.com>

---

mm/vmscan.c | 2 +-  
1 file changed, 1 insertion(+), 1 deletion(-)

```
diff -puN mm/vmscan.c~rss-fix-nodescan mm/vmscan.c
--- linux-2.6.20/mm/vmscan.c~rss-fix-nodescan 2007-05-02 14:47:09.000000000
+0530
+++ linux-2.6.20-balbir/mm/vmscan.c 2007-05-02 14:47:25.000000000 +0530
@@ -1127,7 +1127,7 @@ unsigned long try_to_free_pages_in_conta
    int node;
    struct zone **zones;

- for_each_node(node) {
+ for_each_online_node(node) {
    #ifdef CONFIG_HIGHMEM
        zones = NODE_DATA(node)->node_zonelist[ZONE_HIGHMEM].zones;
        if (do_try_to_free_pages(zones, sc.gfp_mask, &sc))
```

—

--

Warm Regards,  
Balbir Singh  
Linux Technology Center  
IBM, ISTL