
Subject: Re: [PATCH 8/8] Per-container pages reclamation
Posted by [Vaidyanathan Srinivas](#) on Thu, 31 May 2007 08:22:13 GMT
[View Forum Message](#) <> [Reply to Message](#)

Andrew Morton wrote:

> On Wed, 30 May 2007 19:42:26 +0400

> Pavel Emelianov <xemul@openvz.org> wrote:

>
[snip]
>>
>> +#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,
>> + .order = 0, /* in this case we wanted one page only */
>> + .cnt = cnt,
>> + .isolate_pages = isolate_pages_in_container,
>> + };
>> + int node;
>> + struct zone **zones;
>> +
>> + for_each_online_node(node) {
>> +#ifdef CONFIG_HIGHMEM
>> + zones = NODE_DATA(node)->node_zonelists[ZONE_HIGHMEM].zones;
>> + if (do_try_to_free_pages(zones, sc.gfp_mask, &sc))
>> + return 1;
>> +#endif
>> + zones = NODE_DATA(node)->node_zonelists[ZONE_NORMAL].zones;
>> + if (do_try_to_free_pages(zones, sc.gfp_mask, &sc))
>> + return 1;
>> + }
>
> Definitely need to handle ZONE_DMA32 and ZONE_DMA (some architectures put
> all memory into ZONE_DMA (or they used to))

Can we iterate from MAX_NR_ZONES-1 to automatically choose the right
zone for the system?

```
for_each_online_node(node)
for (i = MAX_NR_ZONES - 1; i >= 0; i--) {
    zones = NODE_DATA(node)->node_zonelists[i].zones;
    if (do_try_to_free_pages(zones, sc.gfp_mask, &sc))
        return 1;
```

The above code works on my PPC64 box where all memory is in ZOME_DMA.

```
>> + return 0;
>> +}
>> +#endif
>> +
>> unsigned long try_to_free_pages(struct zone **zones, int order, gfp_t gfp_mask)
>> {
>>   struct scan_control sc = {
> -
> To unsubscribe from this list: send the line "unsubscribe linux-kernel" in
> the body of a message to majordomo@vger.kernel.org
> More majordomo info at http://vger.kernel.org/majordomo-info.html
> Please read the FAQ at http://www.tux.org/lkml/
>
```
