
Subject: Re: [PATCH 8/8] Per-container pages reclamation
Posted by [Balbir Singh](#) on Thu, 31 May 2007 09:22:31 GMT

[View Forum Message](#) <> [Reply to Message](#)

Vaidyanathan Srinivasan wrote:

>
> 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))
>

Won't the node_zonelists[ZONE_NORMAL].zones contain all lower zones (including ZONE_DMA and ZONE_DMA32) from various nodes organized by cost?

I need to go lookup the node_zonelists creation.

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL
