
Subject: Re: [PATCH 8/8] Per-container pages reclamation

Posted by [Balbir Singh](#) on Fri, 01 Jun 2007 09:23:30 GMT

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

Pavel Emelianov wrote:

> Balbir Singh wrote:

>
> [snip]
>
>>>> + 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.
>
>
> OK, I've got it. The node's zonelist indexed with ZONE_XXX contains
> the zone specified and all the zones beyond it. Thus shrinking from
> ZONE_HIGHMEM (ZONE_NORMAL if the latter does not exist) is enough.
>
> Thanks,
> Pavel

Exactly!

--

Warm Regards,
Balbir Singh
Linux Technology Center
IBM, ISTL
