
Subject: Re: [PATCH 7/8] Scanner changes needed to implement per-container scanner

Posted by [Andrew Morton](#) on Wed, 30 May 2007 21:46:47 GMT

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

On Wed, 30 May 2007 19:39:41 +0400

Pavel Emelianov <xemul@openvz.org> wrote:

```
> The core change is that the isolate_lru_pages() call is
> replaced with struct scan_controll->isolate_pages() call.
>
> Other changes include exporting __isolate_lru_page() for
> per-container isolator and handling variable-to-pointer
> changes in try_to_free_pages().
>
> This makes possible to use different isolation routines
> for per-container page reclamation. This will be used by
> the following patch.
>
> ...
>
> +struct rss_container;
> +extern unsigned long try_to_free_pages_in_container(struct rss_container *);
> +int __isolate_lru_page(struct page *page, int mode);
>
> extern unsigned long shrink_all_memory(unsigned long nr_pages);
> extern int vm_swappiness;
> extern int remove_mapping(struct address_space *mapping, struct page *page);
> diff -upr linux-2.6.22-rc2-mm1.orig/mm/vmscan.c linux-2.6.22-rc2-mm1-0/mm/vmscan.c
> --- linux-2.6.22-rc2-mm1.orig/mm/vmscan.c 2007-05-30 12:32:36.000000000 +0400
> +++ linux-2.6.22-rc2-mm1-0/mm/vmscan.c 2007-05-30 16:13:09.000000000 +0400
> @@ -47,6 +47,8 @@
>
> #include "internal.h"
>
> +#include <linux/rss_container.h>
> +
> struct scan_control {
> /* Incremented by the number of inactive pages that were scanned */
> unsigned long nr_scanned;
> @@ -70,6 +72,13 @@ struct scan_control {
> int all_unreclaimable;
>
> int order;
> +
> + struct rss_container *cnt;
```

Can we please have a better name? "cnt" is usually a (poorly-chosen) name for an integer counter. Perhaps "container", or even "rss_container".

```
> + nr_reclaimed += shrink_zones(priority, zones, sc);  
> + if (sc->cnt == NULL)  
> + shrink_slab(sc->nr_scanned, gfp_mask, lru_pages);
```

We don't we shrink slab if called to shrink a container.

This is a fundamental design decision, and a design shortcoming. A full discussion of this is absolutely appropriate to the patch changelog. Please don't just hide stuff like this in the patch and leave people wondering, or ignorant.
