

# Re: [patch 09/10] Remove the SLOB allocator for 2.6.23

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-07/msg03917.html>

---

- *From:* Christoph Lameter <[clameter@xxxxxxx](mailto:clameter@xxxxxxx)>
  - *Date:* Wed, 11 Jul 2007 11:33:52 -0700 (PDT)
- 

On Wed, 11 Jul 2007, Pekka J Enberg wrote:

Hi Christoph,

On Wed, 11 Jul 2007, Christoph Lameter wrote:

Of course you are the maintainer but you only authored a single patch which was the original submission in all the time that SLOB was in the tree. I keep having to clean up the allocator that has--according to Pekka--more memory requirements than SLUB. There is no point in keeping it around anymore it seems.

Well, it was a test setup with UML and busybox and didn't have all the SLOB optimizations Nick mentioned, so we shouldn't draw any definite conclusions from it. I couldn't get 2.6.22-rc6-mm1 to compile so I'll try again after Andrew pushes a new release out.

But you did get 2.6.22 to compile it seems.

Here is the fix against 2.6.22-rc6-mm1 again.

---  
mm/slub.c | 4 ++++  
1 file changed, 4 insertions(+)

Index: linux-2.6.22-rc6-mm1/mm/slub.c

```
=====
--- linux-2.6.22-rc6-mm1.orig/mm/slub.c 2007-07-06 13:28:57.000000000 -0700
+++ linux-2.6.22-rc6-mm1/mm/slub.c 2007-07-06 13:29:01.000000000 -0700
@@ -1868,7 +1868,9 @@ static void init_kmem_cache_node(struct
atomic_long_set(&n->nr_slabs, 0);
spin_lock_init(&n->list_lock);
INIT_LIST_HEAD(&n->partial);
+#ifdef CONFIG_SLUB_DEBUG
```

Re: [patch 09/10] Remove the SLOB allocator for 2.6.23

```
INIT_LIST_HEAD(&n->full);
+#endif
}

#ifdef CONFIG_NUMA
@@ -1898,8 +1900,10 @@ static struct kmem_cache_node * __init e
page->freelist = get_freepointer(kmalloc_caches, n);
page->inuse++;
kmalloc_caches->node[node] = n;
#ifdef CONFIG_SLUB_DEBUG
init_object(kmalloc_caches, n, 1);
init_tracking(kmalloc_caches, n);
#endif
init_kmem_cache_node(n);
atomic_long_inc(&n->nr_slabs);
add_partial(n, page);
-
```

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxxxx  
More majordomo info at <http://vger.kernel.org/majordomo-info.html>  
Please read the FAQ at <http://www.tux.org/lkml/>