

## [PATCH] Re: Kernel Panic – 2.6.23–rc4–mm1 ia64 – was Re: Update: [Automatic] NUMA replicated pagecache ...

---

*Source:* <http://linux.derkeiler.com/Mailing-Lists/Kernel/2007-09/msg03228.html>

---

- *From:* Lee Schermerhorn <[Lee.Schermerhorn@xxxxxx](mailto:Lee.Schermerhorn@xxxxxx)>
  - *Date:* Wed, 12 Sep 2007 15:46:14 -0400
- 

On Wed, 2007-09-12 at 16:41 +0100, Andy Whitcroft wrote:

On Wed, Sep 12, 2007 at 11:09:47AM -0400, Lee Schermerhorn wrote:

Interesting, I don't see a memory controller function in the stack trace, but I'll double check to see if I can find some silly race condition in there.

right. I noticed that after I sent the mail.

Also, config available at:

<http://free.linux.hp.com/~lts/Temp/config-2.6.23-rc4-mm1-gwydyr-nomemcont>

Be interested to know the outcome of any bisect you do. Given its tripping in reclaim.

Problem isolated to memory controller patches. This patch seems to fix this particular problem. I've only run the test for a few minutes with and without memory controller configured, but I did observe reclaim kicking in several times. W/o this patch, system would panic as soon as I entered direct/zone reclaim--less than a minute.

Lee

-----

PATCH 2.6.23-rc4-mm1 Memory Controller: initialize all scan\_controls' isolate\_pages member.

We need to initialize all scan\_controls' isolate\_pages member. Otherwise, shrink\_active\_list() attempts to execute at undefined location.

[PATCH] Re: Kernel Panic – 2.6.23-rc4-mm1 ia64 – was Re: Update: [Automatic] NUMA replicated pagecache ...

Signed-off-by: Lee Schermerhorn <lee.schermerhorn@xxxxxx>

mm/vmscan.c | 2 ++  
1 file changed, 2 insertions(+)

Index: Linux/mm/vmscan.c

```
-----  
--- Linux.orig/mm/vmscan.c 2007-09-10 13:22:21.000000000 -0400  
+++ Linux/mm/vmscan.c 2007-09-12 15:30:27.000000000 -0400  
@@ -1758,6 +1758,7 @@ unsigned long shrink_all_memory(unsigned  
.swap_cluster_max = nr_pages,  
.may_writepage = 1,  
.swappiness = vm_swappiness,  
+ .isolate_pages = isolate_pages_global,  
};  
  
current->reclaim_state = &reclaim_state;  
@@ -1941,6 +1942,7 @@ static int __zone_reclaim(struct zone *z  
SWAP_CLUSTER_MAX),  
.gfp_mask = gfp_mask,  
.swappiness = vm_swappiness,  
+ .isolate_pages = isolate_pages_global,  
};  
unsigned long slab_reclaimable;
```

–  
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/>

[PATCH] Re: Kernel Panic – 2.6.23-rc4-mm1 ia64 – was Re: Update: [Automatic] NUMA replicated pagecache ...