

Re: [PATCH 14/14] page-replace-kswapd-incmin.patch

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-12/msg08594.html>

- *From:* Peter Zijlstra <a.p.zijlstra@xxxxxxxx>
 - *Date:* Sat, 31 Dec 2005 10:40:40 +0100
-

On Fri, 2005-12-30 at 23:15 -0200, Marcelo Tosatti wrote:

> On Fri, Dec 30, 2005 at 11:42:34PM +0100, Peter Zijlstra wrote:

>>

>> From: Nick Piggin <npiggin@xxxxxxx>

>>

>> Explicitly teach kswapd about the incremental min logic instead of just scanning
>> all zones under the first low zone. This should keep more even pressure applied
>> on the zones.

>>

>> The new shrink_zone() logic exposes the very worst side of the current
>> balance_pgdat() function. Without this patch reclaim is limited to ZONE_DMA.

>

> Can you please describe the issue with over protection of DMA zone you experienced?

>

> I'll see if I can reproduce it with Nick's standalone patch on top of vanilla, what
> load was that?

With the mdb bench the following behaviour was observed:

(mem=128M)

- PageCache would fill zone_normal
- PageCache would fill zone_dma
- reclaim starts
- initially things look right
- after a while zone_dma is reclaimed so fast that it frequently gets a full eviction (nr_resident == 0).
- from this point onward zone_normal practically sits idle and zone_dma goes wild with all the action.

--

Peter Zijlstra <a.p.zijlstra@xxxxxxxx>

-

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/>

- **References:**

- ◆ **[PATCH] vm: page-replace and clockpro**
 - ◇ From: Peter Zijlstra
- ◆ **[PATCH 14/14] page-replace-kswapd-incmin.patch**
 - ◇ From: Peter Zijlstra
- ◆ **Re: [PATCH 14/14] page-replace-kswapd-incmin.patch**
 - ◇ From: Marcelo Tosatti

- Prev by Date: **Re: [PATCH] Fix user data corrupted by old value return of sysctl**
- Next by Date: **Re: [PATCH] Fix user data corrupted by old value return of sysctl**
- Previous by thread: **Re: [PATCH 14/14] page-replace-kswapd-incmin.patch**
- Next by thread: **[PATCH 3/9] clockpro-PG test.patch**
- Index(es):
 - ◆ **Date**
 - ◆ **Thread**