
Subject: Re: [PATCH 3/5] Switch caches notification dynamically
Posted by [Pavel Emelianov](#) on Mon, 01 Oct 2007 13:21:58 GMT
[View Forum Message](#) <> [Reply to Message](#)

```
>> diff --git a/mm/slub.c b/mm/slub.c
>> index 31d04a3..e066a0e 100644
>> --- a/mm/slub.c
>> +++ b/mm/slub.c
>> @@ -3779,6 +3932,60 @@ static ssize_t defrag_ratio_store(struct
>> SLAB_ATTR(defrag_ratio);
>> #endif
>>
>> +static ssize_t cache_notify_show(struct kmem_cache *s, char *buf)
>> +{
>> + return sprintf(buf, "%d\n", !(s->flags & SLAB_NOTIFY));
>> +}
>> +
>> +static ssize_t cache_notify_store(struct kmem_cache *s,
>> + const char *buf, size_t length)
>> +{
>> + int err;
>> + int n;
>> +
>> + if ((buf[0] == '1') && !(s->flags & SLAB_NOTIFY)) {
>
> Won't this code break if I pass 10 as input in buf?
```

I don't care about this :) And no /sys/slub/<xxx>/ file does.
It's the problem of a user if he pass some shit and doesn't
see the notifications turned on/off ;)

```
>> + err = slub_on_notify(s);
>> + if (err < 0)
>> + return err;
>> +
>> + s->flags |= SLAB_NOTIFY;
>> +
>> + flush_all(s);
>
> Sounds like an expensive operation, can we add a comment
> to indicate the same. Also, in the documentation, could
> we warn the user that turning notifications on/off
> can be expensive?
```

Sure.

```
>> + for_each_node_state(n, N_NORMAL_MEMORY) {
>> + struct kmem_cache_node *node;
```

```

>> + struct page *pg;
>> +
>> + node = get_node(s, n);
>> + spin_lock_irq(&node->list_lock);
>> + list_for_each_entry(pg, &node->partial, lru)
>> +   SetSlabDebug(pg);
>> + spin_unlock_irq(&node->list_lock);
>> + }
>> + return length;
>> + }
>> +
>> + if ((buf[0] == '0') && (s->flags & SLAB_NOTIFY)) {
>
> What happens if buf is 01?
>
>> + /*
>> +  * TODO: make the notifications-off work
>> +  */
>> + if (any_slab_objects(s))
>> +   return -EBUSY;
>> +
>> + s->flags &= ~SLAB_NOTIFY;
>> + err = slub_off_notify(s);
>> + if (err < 0) {
>> +   s->flags |= SLAB_NOTIFY;
>> +   return err;
>> + }
>> +
>> + return length;
>> + }
>> +
>> + return -EINVAL;
>> +}
>> +
>> +SLAB_ATTR(cache_notify);
>> +
>> static struct attribute * slab_attrs[] = {
>>   &slab_size_attr.attr,
>>   &object_size_attr.attr,
>>   @@ -3809,6 +4016,7 @@ static struct attribute * slab_attrs[] =
>>   #ifdef CONFIG_NUMA
>>   &defrag_ratio_attr.attr,
>>   #endif
>>   &cache_notify_attr.attr,
>>   NULL
>> };
>>
>>

```

>
>