Subject: Re: [PATCH 2/5] Generic notifiers for SLUB events Posted by Pavel Emelianov on Mon, 01 Oct 2007 13:07:44 GMT

View Forum Message <> Reply to Message

```
Balbir Singh wrote:
> Pavel Emelyanov wrote:
>> If the user wants more than one listener for SLUB event,
>> it can register them all as notifier_block-s so all of
>> them will be notified.
>>
>> This costs us about 10% of performance loss, in comparison
>> with static linking.
>>
>> The selected method of notification is srcu notifier blocks.
>> This is selected because the "call" path, i.e. when the
>> notification is done, is lockless and at the same time the
>> handlers can sleep. Neither regular nor atomic notifiers
>> provide such facilities.
>>
>> Signed-off-by: Pavel Emelyanov <xemul@openvz.org>
>>
>> ---
>>
>> diff --git a/init/Kconfig b/init/Kconfig
>> index 684ccfb..e9acc29 100644
>> --- a/init/Kconfig
>> +++ b/init/Kconfig
>> @ @ -593,6 +599,16 @ @ config SLUB DEBUG
     SLUB sysfs support. /sys/slab will not exist and there will be
>>
     no support for cache validation etc.
>>
>> +config SLUB_NOTIFY
>> + default y
> Should the default be on? Shouldn't it depend on KMEM?
Well... I think that is should be N by default and has
noting to do with the KMEM:) Thanks for noticing.
>> + bool "Enable SLUB events generic notification"
>> + depends on SLUB
>> + help
>> + When Y, this option enables generic notifications of some major
>> + slub events. However, if you do know that there will be the
>> + only listener for them, you may say N here, so that callbacks
>> + will be called directly.
>> +
```

>> -

```
>> -
>> +#ifdef CONFIG_SLUB_NOTIFY
>> + srcu_init_notifier_head(&slub_nb);
>
> Can we get rid of the #ifdef CONFIG_SLUB_NOTIFY?

I don't think this is really useful in the __init code :)
>> +#endif
>> printk(KERN_INFO "SLUB: Genslabs=%d, HWalign=%d, Order=%d-%d, MinObjects=%d,"
>> " CPUs=%d, Nodes=%d\n",
>> caches, cache_line_size(),
>> >
> >
```