

[patch] preempt-trace.patch

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-10/7077.html>

From: Ingo Molnar (mingo_at_elte.hu)

Date: 10/31/05

Date: Mon, 31 Oct 2005 22:42:02 +0100

To: Andrew Morton <akpm@osdl.org>

* Andrew Morton <akpm@osdl.org> wrote:

> *Well I can't find it. Ingo, didn't you have a debug patch which would
> help us identify where this atomic section started?*

yes. Find below the preempt-trace.patch, updated and merged to 2.6.14.

Rafael, just enable CONFIG_DEBUG_PREEMPT and repeat the reproducer.

Works fine on x86 and x64, and should not break other architectures. Has been included in the -rt tree for quite some time. It is certainly a useful debugging feature. Could we add this to -mm, with an focus on getting this upstream eventually?

Ingo

this implements the "non-preemptible section trace" feature, which prints out a "critical section nesting" trace after stackdumps.

sample output on x86:

Call Trace:

```
[<c0103db1>] show_stack+0x7a/0x90  
[<c0103f36>] show_registers+0x156/0x1ce  
[<c010412e>] die+0xe8/0x172  
[<c010422e>] do_trap+0x76/0xa3  
[<c01044fe>] do_invalid_op+0xa3/0xad  
[<c01039ef>] error_code+0x4f/0x54  
[<c0120be9>] test+0x8/0xa  
[<c0120c41>] sys_gettimeofday+0x56/0x74  
[<c0102eeb>] sysenter_past_esp+0x54/0x75
```

```
-----  
| preempt count: 00000004 |  
| 4 levels deep critical section nesting:  
-----
```

Linux-Kernel: [patch] preempt-trace.patch

```
.. [<c0120bbe>] .... test3+0xd/0xf
.....[<c0120bc8>] .. ( <= test2+0x8/0x21)
.. [<c0120bbe>] .... test3+0xd/0xf
.....[<c0120bcd>] .. ( <= test2+0xd/0x21)
.. [<c0120bd7>] .... test2+0x17/0x21
.....[<c0120be9>] .. ( <= test+0x8/0xa)
.. [<c010407f>] .... die+0x39/0x172
.....[<c010422e>] .. ( <= do_trap+0x76/0xa3)
```

sample output on x64:

```
Call Trace:<ffffffff8012373c>{sys32_gettimeofday+60} <ffffffff80122af3>{cstar_do_call+27}
```

```
-----
| preempt count: 00000002 ]
| 2 level deep critical section nesting:
-----
```

```
.. [<ffffffff8012372d>] .... sys32_gettimeofday+0x2d/0x110
.....[<ffffffff80122af3>] .. ( <= cstar_do_call+0x1b/0x60)
.. [<ffffffff80123737>] .... sys32_gettimeofday+0x37/0x110
.....[<ffffffff80122af3>] .. ( <= cstar_do_call+0x1b/0x60)
```

this feature