

Re: [crash, bisected] Re: [PATCH 3/4] x86_64: Fold pda into per cpu area

Re: [crash, bisected] Re: [PATCH 3/4] x86_64: Fold pda into per cpu area

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2008-07/msg11313.html>

- *From:* Mike Travis <travis@xxxxxxx>
 - *Date:* Fri, 25 Jul 2008 14:05:11 -0700
-

Ok, I'll just post what I have now (compiles and boots cleanly)... and then we can discuss these more extensively.

Thanks,
Mike

Jeremy Fitzhardinge wrote:

Mike Travis wrote:

Is this for the boot cpu (0), or for all cpus? For the boot cpu, I have this now in arch/x86/kernel/setup_percpu.c:

```
+#ifdef CONFIG_HAVE_ZERO_BASED_PER_CPU
+
+/* Initialize percpu offset for boot cpu (0) */
+unsigned long __per_cpu_offset[NR_CPUS] __read_mostly = {
+ [0] = (unsigned long) __per_cpu_load
+};
+#else
+unsigned long __per_cpu_offset[NR_CPUS] __read_mostly;
+#endif
```

So this should apply as well to the xen startup?

If it's just a static initialization, then it should be fine. But some equivalent of your head_64.S changes are needed to actually set things up?

xen_cpu_up() needs to do whatever initialization needed for a new cpu's percpu area (presumably whatever do_boot_cpu() does).

Re: [crash, bisected] Re: [PATCH 3/4] x86_64: Fold pda into per cpu area

Does the startup include executing
arch/x86/kernel/head_64.S:startup_64() ?
I see arch/x86/xen/xen-head.S:startup_xen() so I'm guessing not?

No, it doesn't. It bypasses all that startup code. Aside from the few instructions in xen-head.S, xen_start_kernel() is the first thing to get run.

But when bringing up a secondary cpu, where does the new percpu memory actually get allocated?

For the real startup, I do the following two things. But I'm not comfortable enough with xen to think I'll get it right putting this in xen-head.S.

Yes, it needn't be in the asm code. I'll work out what to do. Looks like I just need to do an appropriate wrmsr(MSR_GS_BASE,).

J

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@xxxxxxxxxxxxxxxxx
More majordomo info at <http://vger.kernel.org/majordomo-info.html>
Please read the FAQ at <http://www.tux.org/lkml/>