

Re: [Fastboot] [RFC] [PATCH 2/2] kdump: cciss driver initialization issue fix

Re: [Fastboot] [RFC] [PATCH 2/2] kdump: cciss driver initialization issue fix

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2006-06/msg08364.html>

- *From:* Vivek Goyal <vgoyal@xxxxxxxxxx>
 - *Date:* Mon, 26 Jun 2006 15:36:06 -0400
-

On Mon, Jun 26, 2006 at 01:51:43PM -0500, Miller, Mike (OS Dev) wrote:

[..]

Subject: Re: [Fastboot] [RFC] [PATCH 2/2] kdump: cciss driver initialization issue fix

"Miller, Mike (OS Dev)" <Mike.Miller@xxxxxx> writes:

Thanks Eric, that helps me understand. Section 8.2.2 of the open cciss

spec supports a reset message. Target 0x00 is the controller. We could

add this to the init routine to ensure the board is made sane again but this would drastically increase init time under normal

circumstances.

Where does the init time penalty come from? How large is the init penalty? I suspect it is from waiting for the scsi disks to spin up. But I am just guessing in the dark.

The penalty is in the firmware and self-test operations.

And I suspect this is a hard reset, also. Not sure if that would negatively impact kdump. If there were some condition we could test

Re: [Fastboot] [RFC] [PATCH 2/2] kdump: cciss driver initialization issue fix

against and perform the reset when that condition is met it

would not

impact 99.9% of users.

I am wondering if it is possible to look at the controller and see if it is in a bad state, (i.e. in some state besides just coming out of reset) and if so issue a reset. If this really is a long operation that would be the ideal way to handle it.

It's not really in a bad state at this time, is it? Maybe some commands hanging around.

That's true. Only some commands are hanging around and currently only a warning is printed, if CONFIG_CISS_SCSI_TAPE is not enabled. Otherwise, driver thinks that there is no reasons for me receiving a command completion message for a command which I never issued and it runs in to the BUG().

```
/* This will need changing for direct lookup completions */  
if (complete != c->busaddr) {  
    if (add_sendcmd_reject(cmd, ctlr, complete) != 0) {  
        BUG(); /* we are pretty much hosed if we get here. */  
    }  
    continue;
```

If the amount of time is really user noticeable and testing for it is impossible then it is probably time to talk kernel command line options.

I was informed of the crashboot command line parameter. I can implement that as a test.

Now as per Andrew's suggestion, I have changed crashboot to "reset_devices".

Please find attached the patch. Now reset functionality of cciss driver needs to be implemented.

Thanks
Vivek

Re: [Fastboot] [RFC] [PATCH 2/2] kdump: cciss driver initialization issue fix

- o Introduce "reset_devices" command line option.
- o Resetting the devices during driver initialization can be a costly operation in terms of time (especially scsi devices). This option can be used by drivers to know that user forcibly wants the devices to be reset during initialization.
- o This option can be useful while kernel is booting in unreliable environment. For ex. during kdump boot where devices are in unknown random state and BIOS execution has been skipped.

Signed-off-by: Vivek Goyal <vgoyal@xxxxxxxxxx>

```
linux-2.6.17-1M-vivek/Documentation/kernel-parameters.txt | 3 ++
linux-2.6.17-1M-vivek/include/linux/init.h | 1
linux-2.6.17-1M-vivek/init/main.c | 20 ++++++
3 files changed, 24 insertions(+)
```

```
diff -puN init/main.c~add-reset-devices-command-line-option init/main.c
--- linux-2.6.17-1M/init/main.c~add-reset-devices-command-line-option 2006-06-26
13:53:51.000000000 -0400
+++ linux-2.6.17-1M-vivek/init/main.c 2006-06-26 15:06:12.000000000 -0400
@@ -125,6 +125,18 @@ static char *ramdisk_execute_command;
static unsigned int max_cpus = NR_CPUS;

/*
+ * If set, this is an indication to the drivers that reset the underlying
+ * device before going ahead with the initialization otherwise driver might
+ * rely on the BIOS and skip the reset operation.
+ *
+ * This is useful if kernel is booting in an unreliable environment.
+ * For ex. kdump situation where previous kernel has crashed, BIOS has been
+ * skipped and devices will be in unknown state.
+ */
+unsigned int reset_devices;
+EXPORT_SYMBOL(reset_devices);
+
+/*
* Setup routine for controlling SMP activation
*
* Command-line option of "nosmp" or "maxcpus=0" will disable SMP
@@ -150,6 +162,14 @@ static int __init maxcpus(char *str)

__setup("maxcpus=", maxcpus);

+static int __init set_reset_devices(char *str)
+{
```

Re: [Fastboot] [RFC] [PATCH 2/2] kdump: cciss driver initialization issue fix

```
+ reset_devices = 1;
+ return 1;
+}
+
+__setup("reset_devices", set_reset_devices);
+
static char * argv_init[MAX_INIT_ARGS+2] = { "init", NULL, };
char * envp_init[MAX_INIT_ENVS+2] = { "HOME=/", "TERM=linux", NULL, };
static const char *panic_later, *panic_param;
diff -puN include/linux/init.h~add-reset-devices-command-line-option include/linux/init.h
--- linux-2.6.17-1M/include/linux/init.h~add-reset-devices-command-line-option 2006-06-26
14:44:22.000000000 -0400
+++ linux-2.6.17-1M-vivek/include/linux/init.h 2006-06-26 14:44:51.000000000 -0400
@@ -69,6 +69,7 @@ extern initcall_t __security_initcall_st

/* Defined in init/main.c */
extern char saved_command_line[];
+extern unsigned int reset_devices;

/* used by init/main.c */
extern void setup_arch(char **);
diff -puN Documentation/kernel-parameters.txt~add-reset-devices-command-line-option
Documentation/kernel-parameters.txt
--- linux-2.6.17-1M/Documentation/kernel-parameters.txt~add-reset-devices-command-line-option
2006-06-26 14:45:01.000000000 -0400
+++ linux-2.6.17-1M-vivek/Documentation/kernel-parameters.txt 2006-06-26 14:47:20.000000000 -0400
@@ -1340,6 +1340,9 @@ running once the system is up.
```

reserve= [KNL,BUGS] Force the kernel to ignore some iomem area

+ reset_devices [KNL] Force drivers to reset the underlying device
+ during initialization.

+
resume= [SWSUSP]
Specify the partition device for software suspend

-
-

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/>