

Mark,

Salyzyn, Mark wrote:

- > Vasily, it will necessarily be up to you as to whether you switch to
- > dpt_i2o to get the hardening you require today, or work out a deal with
- > Markus to add timeout/reset functionality to the i2o driver.

Of course, you are right. Currently our customers have had 2 alternatives:

- be tolerate to these hangs
- if they can't bear it -- replace i2o hardware

Therefore first at all I'm going to add third possible alternative, dpt_i2o driver.

Mark, could you please send me latest version of your driver directly? Or can I probably take it from mainstream?

The next task is help Markus in i2o error/reset handler implementation.

- > My recommendations for the i2o driver reset procedure is to use a
- > rolling timeout, every new command completion resets the global timer.
- > This will allow starved or long commands to process. Once the timer hits
- > 3 minutes for RAID (Block or SCSI) targets that have multiple
- > inheritances, 30 seconds for SCSI DASD targets, or some insmod tunable,
- > it resets the adapter. I recommend that when we hit ten seconds, or some
- > insmod tunable, that we call a card specific health check routine. I do
- > not recommend health check polling because we have noticed a reduction
- > in Adapter performance in some systems and generic i2o cards would
- > require a command to check, so that is why I tie it to the ten seconds
- > past last completion. For the DPT/Adaptec series of adapters, it checks
- > the BlinkLED status (code fragment in dpt_i2o driver at
- > adpt_read_blink_led), and if set, immediately record the fact and resets
- > the adapter. For cards other than the DPT/Adaptec series, I recommend a
- > short timeout Get Status request to see if the Firmware is in a run
- > state and is responsive to this simple command. The reset code will need
- > to retry all commands itself, I do not believe the block system has an
- > error status that can be used for it to retry the commands. If the Reset
- > loop in the reset adapter code is unresponsive, then the known targets
- > need to be placed offline.

Sorry, I do not have your big experience in scsi and do not know nothing in i2o. However are you sure than 3 min is enough for timeout? As far as I know some scsi commands (for example rewind on tapes) can last during a very long time.

Also I have some other questions but currently I'm not fell that I'm ready for

this discussion.

Thank you,
Vasily Averin

SWsoft Virtuozzo/OpenVZ Linux kernel team
