Document Number	EDCS-1234848
Based on Template	**EDCS-1234848**
Created By	**D. Francheski**

IPMI Firmware Upgrade

Details the process of upgrading CDE250 IPMI firmware in the field

Reviewers

Department	Name/Title
Development Engineering	Bilal Raja
Development Engineering	Scott Parry
DevTest Engineering	

The departments and/or individuals listed above should be notified in advance and given a sufficient time period to review this document. The Project Team determines requirements for approval according to the scope of the project.

Modification History

Revision	Date	Originator	Comments
1	3/6/2013	Scott Parry / David Francheski	Initial Draft
1.1	4/9/2013	Scott Parry	Added item 3 to the Background section
1.2	4/24/2013	Scott Parry	Updated section 2.1
1.3	5/1/2013	Scott Parry	Updated USB FreeDOS Boot Stick section

Contents

1 Background

Current CDE250 systems utilize an Intelligent Platform Management Interface (IPMI) infrastructure that monitors and manages the health of the system. This is implemented using embedded IPMI firmware running within the Baseboard Management Controller (BMC) system chip on the motherboard. The current version of IPMI firmware used on the CDE250 is v1.33.

There are currently three (3) known problems associated with the IPMI / BMC subsystem and the IPMI FW version 1.33.

1) CPU usage of [kipmi0] process periodically consumes ~ 100% utilization.
   This was recently reported in the following SR customer case:

   SR-624629217 - CDRENS05 / 2.6.3b45 / new CDE250-2S6 CPU usage

   An IBM linux kernel article described this issue as follows:

   " The increased CPU usage of [kipmi] is normal. The hardware interface (BMC) is not interrupt driven, so the driver must poll the device for status and messages. It is this polling that is showing up as a busy CPU. The [kipmi] kernel thread has its priority intentionally lowered so that it does not interfere with other processes on the system. Even when it is polling in its tightest loop (usually when it thinks the BMC has active events it needs to handle), it will give up the CPU to any process that wants it. "

   Due to a defect in earlier versions of IPMI firmware (including v1.33), this polling mechanism was inefficient, often times resulting in high CPU usage. The Cisco supplier of the CDE250 systems has released a new version of IPMI firmware (v2.05), that resolves this CPU utilization issue. We have validated this new firmware internally at Cisco and confirmed it acceptably minimizes [kipmi] CPU usage. IPMI v2.05 firmware will be shipped on all new CDE250 systems out of manufacturing. Please follow the procedure described in Chapter 2 to upgrade the IPMI firmware to v2.05.

2) IPMI Watchdog error messages in kernel system log.
   The following IPMI messages periodically occur in the kernel system log (syslog.txt):

   ```
   %SE-SYS-3-900000: IPMI Watchdog: response: Error ff on cmd 22
   %SE-SYSMON-3-445000: Error in sdt read (hang), restart sdt.
   ```

   The following customer SR cases have been raised against this issue:

   SR-623932327 - Failure of the Cisco CDE-250 Box.

SR-624697161 - CDE250 Load testing threshold with smooth streaming
SR-623663571: Rootfs partition out of free inodes

The later SR (**623663571)** was resolved by fixing the following Cisco defect (resolved in Release 3.1.2 and beyond):

CSCud57383: sysmon robustness against sdt hangs

While the software fix for this defect does not resolve the core underlying IPMI error messages (more on this later), it does significantly reduce the frequency of their occurrence by slowing down the execution interval of the Supero Doctor (**sdt**) hardware health monitoring utility from 5 minutes to 30 minutes (i.e., **sdt** is now run every 30 minutes). **sdt** calls into the IPMI firmware layer to retrieve hardware health information (primarily power-supply status / fan speeds, and various motherboard temperatures / voltages).

[Given that most hardware health characteristics (e.g., fan speeds, temperatures, voltages) normally change at a slow rate, it is desirable to reduce the load on the IPMI subsystem by polling hardware health status only when necessary (i.e., every 30 minutes). If there is an abrupt change in hardware health, the CDE250 motherboard itself has failsafe circuitry to alarm and (if necessary) automatically power-off the system.]

This software fix also properly cleans up / preserves file system resources (in the event the **sdt** utility hangs indefinitely due to the above IPMI issue), thereby preventing inodes from being depleted from the rootfs partition.

The root-cause of the underlying IPMI error messages is being actively researched by both Cisco and the CDE250 system vendor. Upgrading the IPMI firmware to v2.05 will serve to mitigate the risk of this IPMI Watchdog condition occurring, while at the same time improving the stability of the IPMI / BMC subsystem.

At the present time, the only known work-around for this issue is to power-cycle the system. This would suggest that it might be an issue with the BMC controller hardware itself. By cold-restarting the system, any BMC lockup / instability issues are properly cleared upon subsequent power-up. The recommended way to reliably know if the power-cycle / reboot has cleared this issue is to execute a "**show tech**" CLI command. If the command does not hang at the "**system health**" section (this is the output from **sdt**), then the system booted cleanly and should be free of any subsequent IPMI Watchdog issues.

3) POST Boot Hang and POST after system reboot.

After the CDE250 has been running for 24+ hours and a reboot command is issued the system will hang during POST with 5 short and 1 long beep. This pattern is a POST boot error code:

5 short beeps + 1 long beep, Memory error, No memory detected in the system

The only known workaround once the system is in this POST Boot hung state is to remove the power source from the power supplies, wait a few seconds for the power to settle out and then restore power to the system.

The root-cause is the SMBus receiving multiple requests simultaneously which can lock the SMBus, which in turn can lock the memory bus from being accessed by the system during a reboot.

Root cause from the Release notes for IPMI FW v2.05:

Synopsis: Improved SMBus performance.

Description: Reduce system resource usage and decrease the chance of SMBus conflict. If the different sources access the SMBus (I2C) interface at the same time, it could cause system to be hung since the SMBus cannot handle multiple simultaneous requests. SMBus conflicts always produce system a hung on reboot.

The following customer SR cases have been raised against this issue:

SR-624384961 – CDE250-K9-only pull out power supply, than can boot up

Application Lifecycle Management
Defect ID : 226
CDE-250 Vault did not recover from Init-6 reboot
CDS Release 3.2.1-es3 B2
Device : CDE-250.When provisioning SNMP on the new vault the services were not starting so an Init-6 was attempted.

ECO#: E113357 – This ECO is already in place to update the IPMI FW v1.33 with IPMI FW v2.05 in manufacturing

2 IPMI Firmware Upgrade Process

Please follow the process below in order to upgrade the IPMI motherboard firmware to v2.05.

1) FreeDOS USB Boot stick creation procedure

a) Insert a blank USB memory stick into your Windows based PC
b) Download and save the "rufus" USB boot device creator from the following URL: http://rufus.akeo.ie
c) Start rufus by opening the "rufus_vx.x.x.exe" file. (note: this document was created using rufus v1.2.0.183)
d) Insert a USB stick to be formatted with FreeDOS
e) Select the "Device" to use as the FreeDOS boot device
f) File system: "FAT32"
g) Enter a "New volume label", i.e. IPMIUP
h) Select "Quick Format"
i) Select Create a bootable disk using: "FreeDOS"
j) Click on "Start", the Boot stick format should proceed

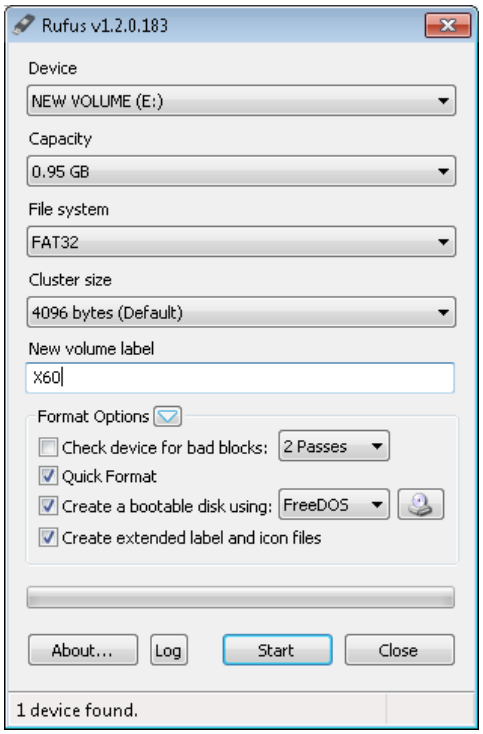

Figure 1: Screenshot of the Rufus utility

2) After the formatting process is complete, you will need to copy the IPMI firmware upgrade utility and associated firmware image onto the USB DOS stick. These are contained in the **ipmiUpgrade.zip p**ackage. The download location for these files is referenced in the Field Notice.

   a) Unzip **ipmiUpgrade.zip** to a temporary folder.
   b) Copy the **ipmiUpgrade/** directory to the root of the USB DOS stick.

3) Remove the USB DOS stick, and insert into any of the available USB ports on the back of the CDE250.

4) Connect a VGA video monitor and keyboard to the back of the CDE250.

5) Boot off the USB DOS stick.
   To accomplish this, simply reload (reboot) the CDE250 from the CLI; it should automatically boot off the USB DOS stick.

   (All CDE250 systems are pre-configured at the factory to first boot off external USB devices, followed by the internal compact flash. If the system doesn't automatically boot off the USB DOS stick, please verify the boot device priority in the BIOS.)

6) After the system boots into a DOS prompt, cd to the **ipmiUpgrade/** directory (example shows a DOS environment that only supports short file names)

   C:\> CD IPMIUP~1

7) To initiate the IPMI firmware upgrade, execute the **IPMIUP.BAT** DOS batch file.

   C:\IPMIUP~1> IPMIUP

   The following will be displayed:

   Update IPMI FW to v2.05

   Current and proposed IPMI BIOS update versions
   … etc …

   Updating the IPMI BIOS to v2.05
   … etc …

   The entire firmware upgrade should take approximately 10 minutes to complete.

Please wait for a positive acknowledge from the utility prior to rebooting.

Please DO NOT power-cycle and/or interrupt the firmware upgrade process. Doing so may result in the system being placed into a non-recoverable state.

8) After the IPMI firmware upgrade completes, you must power-cycle the system. Of course, be sure to first remove the USB DOS stick prior to power-cycling.

2.1 IPMI Firmware Validation

After the IPMI firmware upgrade has been completed and the system power-cycled, validation of a successful upgrade should be performed.

2.1.1 Validation using the IPMI port

Login to IPMI GUI and the login page / System Information Page shows the IPMI firmware Revision.
The Firmware Revision should be "2.05" as shown in the image below:

2.1.2 Validation on CDS-TV

Verify the IPMI FW using the following command:
CDE250_Streamer# cat /sys/devices/platform/ipmi_bmc.aabb.32/firmware_revision
This should return "2.5"

Note: Linux will usually strip the "0" from the 2.05 string.

Execute the "**sdt**" CLI command. Verify this command does not hang (example output is shown below):

```
CDE250_Streamer# sdt

********************************************************************************
 Supero Doctor II - Linux version 2.79(110504)
 Copyright(c) 1993-2010 by Super Micro Computer, Inc. http://supermicro.com/
********************************************************************************
Monitored Item            High Limit  Low Limit     Status
-------------------------------------------------------------------------
Fan1 Fan Speed                             715       5625
Fan2 Fan Speed                             715       5625
Fan7 Fan Speed                             715       5625
CPU1 Vcore Voltage               1.49       0.60       0.98
CPU2 Vcore Voltage               1.49       0.60       0.98
CPU1 VTT Voltage                 1.34       0.92       1.12
CPU2 VTT Voltage                 1.65       0.92       1.13
CPU1 DIMM Voltage                1.65       1.20       1.52
CPU2 DIMM Voltage                1.65       1.20       1.52
+1.5V Voltage                    1.65       1.35       1.46
+1.8V Voltage                    1.98       1.62       1.80
+5V Voltage                      5.50       4.51       4.90
+12V Voltage                    13.19      10.80      12.23
+1.1V Voltage                    1.21       0.98       1.10
+3.3V Voltage                    3.65       2.95       3.29
+3.3Vsb Voltage                  3.65       2.95       3.26
VBAT Voltage                     3.65       2.95       3.02
CPU1 Temperature               95/203                   Low
CPU2 Temperature               95/203                   Low
System Temperature             75/167                 31/87
Chassis Intrusion                                      Good
Power Supply Failure                                   Good
Power1 Supply Failure                                   N/A   Warning!
```

```
Power2 Supply Failure                                    N/A   Warning!
---------------------------------------- Wed Mar  6 16:31:31 2013
```

Note: if sdt returns the following error:

```
**********************************************************************
 SuperDoctor II - Linux version 2.88(120726)
 Copyright(c) 1993-2012 by Super Micro Computer, Inc. http://supermicro.com/
**********************************************************************
Error: no i2c device files found. Please use mkdev.sh to create them.
```

Run the following command: /usr/sbin/enable_ipmi.sh
Then run the "sdt" CLI command again.

2.1.3 Validation on VDS-IS (Release 3.2)

For Release 3.2 (and beyond), simply execute the "**show hardware health**" CLI command. Verify this command does not hang, and that the IPMI firmware version indicates "**2.5**" (example output is shown below):

```
DD13-2S10-2#sh hardware health

 ---------------- System Hardware Health ------------------

**************************************************************************
  Supero Doctor II - Linux version 2.79(110504)
  Copyright(c) 1993-2010 by Super Micro Computer, Inc. http://supermicro.com/
**************************************************************************
Monitored Item            High Limit  Low Limit     Status
-------------------------------------------------------------------------
Fan1 Fan Speed                             715       5625
Fan2 Fan Speed                             715       5273
Fan7 Fan Speed                             715       5273
Power1 Fan 1 Speed                        4007       7099
Power1 Fan 2 Speed                        4007      10992
CPU1 Vcore Voltage                1.49     0.60       1.04
CPU2 Vcore Voltage                1.49     0.60       0.98
CPU1 VTT Voltage                  1.34     0.92       1.14
CPU2 VTT Voltage                  1.65     0.92       1.12
CPU1 DIMM Voltage                 1.65     1.20       1.52
CPU2 DIMM Voltage                 1.65     1.20       1.52
+1.5V Voltage                     1.65     1.35       1.46
+1.8V Voltage                     1.98     1.62       1.80
+5V Voltage                       5.50     4.51       5.09
+12V Voltage                     13.19    10.80      12.07
+1.1V Voltage                     1.21     0.98       1.10
+3.3V Voltage                     3.65     2.95       3.29
```

```
+3.3Vsb Voltage                  3.65       2.95       3.24
VBAT Voltage                     3.65       2.95       2.86   Warning!
CPU1 Temperature                95/203                  Low
CPU2 Temperature                95/203                  Low
System Temperature              75/167                30/86
Power1 Temperature              75/167                34/93
Chassis Intrusion                                      Good
Power Supply Failure                                   Good
Power1 Supply Failure                                  Good
Power2 Supply Failure                                  Bad    Warning!
---------------------------------------- Thu Jan 17 06:37:09 2013
Power1 Information: Type = AC Power, Firmware Rev = 6.0
Power2 Information: Type = Unknown
IPMI Information: Firmware Rev = 2.5
```

2.1.4 Validation on VDS-IS (Pre-Release 3.2)

For systems running VDS-IS software prior to Release 3.2, the following signed script should be executed to verify the IPMI firmware version. The download location for this file is referenced in the Field Notice.

```
#script execute showipmifw.sh.signed
IPMI Firmware Revision 2.5

Script showipmifw.sh.signed exited with return code 0
```

In addition, please issue a "**show tech**" CLI command. As stated earlier, if the command does not hang at the "**system health**" section (output from **sdt**), then the system booted cleanly and should be free of any subsequent IPMI Watchdog issues.

2.1.5 Validation on VQE

Verify the IPMI FW using the following command at the console:

root@]# ipmiutil sel show

This command should return the following value for the "BMC version"

```
ipmiutil ver 2.51
showsel: version 2.51
-- BMC version 2.5, IPMI version 2.0
SEL Ver 51 Support 3, Used = 1 records, Free space = 574 records
RecId Date/Time_______ Source_ Evt_Type SensNum Evt_detail - Trig [Evt_data]
0001 04/19/13 02:58:43 BMC   08 Power Supply #1a is OK    ef [01 ff ff]
showsel completed successfully
```

Note: Linux will usually strip the "0" from the 2.05 string.