



Logging

Table of contents

RShim Logging

IPMI Logging in UEFI

SEL Record Format

Possible SEL Field Values

Event Definitions

Reading IPMI SEL Log Messages

ACPI BERT Logging

RShim Logging

RShim logging uses an internal 1KB HW buffer to track booting progress and record important messages. It is written by the NVIDIA® BlueField® networking platform's (DPU or SuperNIC) Arm cores and is displayed by the RShim driver from the USB/PCIe host machine. Starting in release 2.5.0, ATF has been enhanced to support the RShim logging.

The RShim log messages can be displayed described in the following:

1. Check the DISPLAY_LEVEL level in file /dev/rshim0/misc.

```
# cat /dev/rshim0/misc
DISPLAY_LEVEL 0 (0:basic, 1:advanced, 2:log)
...
```

2. Set DISPLAY_LEVEL to 2.

```
# echo "DISPLAY_LEVEL 2" > /dev/rshim0/misc
```

3. Log messages are displayed in the misc file.

```
# cat /dev/rshim0/misc
...
-----
      Log Messages
-----
INFO[BL2]: start
INFO[BL2]: no DDR on MSS0
INFO[BL2]: calc DDR freq (clk_ref 53836948)
INFO[BL2]: DDR POST passed
INFO[BL2]: UEFI loaded
INFO[BL31]: start
INFO[BL31]: runtime
INFO[UEFI]: eMMC init
INFO[UEFI]: eMMC probed
```

INFO[UEFI]: PCIe enum start
 INFO[UEFI]: PCIe enum end

i Info

This is an example output for BlueField-2.

The following table details the ATF/UEFI messages for BlueField-2 and BlueField-3:

Message	Explanation	Action
INFO[BL2]: start	BL2 started	Informational
INFO[BL2]: no DDR on MSS<N>	DDR is not detected on memory controller <N>	Informational (depends on device)
INFO[BL2]: calc DDR freq (clk_ref 156M, clk xxx)	DDR frequency is calculated based on reference clock 156M	Informational
INFO[BL2]: calc DDR freq (clk_ref 100M, clk xxx)	DDR frequency is calculated based on reference clock 100M	Informational
INFO[BL2]: calc DDR freq (clk_ref xxxx)	DDR frequency is calculated based on reference clock xxxx	Informational
INFO[BL2]: DDR POST passed	BL2 DDR training passed	Informational
INFO[BL2]: UEFI loaded	UEFI image is loaded successfully in BL2	Informational
ERR[BL2]: DDR init fail on MSS<N>	DDR initialization failed on memory controller <N>	Informational (depends on device)
ERR[BL2]: image <N> bad CRC	Image with ID <N> is corrupted which will cause hang	Error message. Reset the device and retry. If problem persists, use

Message	Explanation	Action
		a different image to retry it.
ERR[BL2]: DDR BIST failed	DDR BIST failed	Need to retry. Check the ATF booting message whether the detected OPN is correct or not, or whether it is supported by this image. If still fails, contact NVIDIA Support.
ERR[BL2]: DDR BIST Zero Mem failed	DDR BIST failed in the zero-memory operation	Power-cycle and retry. If the problem persists, contact your NVIDIA FAE.
WARN[BL2]: DDR frequency unsupported	DDR training is programmed with unsupported parameters	Check whether official FW is being used. If the problem persists, contact your NVIDIA FAE.
WARN[BL2]: DDR min-sys(unknown)	System type cannot be determined and boot as a minimal system	Check whether the OPN or PSID is supported. If the problem persists, contact your NVIDIA FAE.
WARN[BL2]: DDR min-sys(misconf)	System type misconfigured and boot as a minimal system	Check whether the OPN or PSID is supported. If the problem persists, contact your NVIDIA FAE.
Exception(BL2): syndrome = xxxxxxxx ...	Exception in BL2 with syndrome code and register dump. System hung.	Capture the log, analyze the cause, and report to FAE if needed
PANIC(BL2): PC = xxx ...	Panic in BL2 with register dump. System will hung.	Capture the log, analyze the cause, and report to FAE if needed
ERR[BL2]: load/auth failed	Failed to load image (non-existent/corrupted), or image authentication failed when secure boot is enabled	Try again with the correct and properly signed image
INFO[BL31]: start	BL31 started	Informational
INFO[BL31]: runtime	BL31 enters the runtime state. This is the latest BL31 message in normal booting process.	Informational

Message	Explanation	Action
Exception(BL31): syndrome = xxxxxxxx cptr_el3 xx daif xx ...	Exception in BL31 with syndrome code and register dump. System hung.	Capture the log, analyze the cause, and report to FAE if needed
PANIC(BL31): PC = xxx cptr_el3 xxx daif xxx ...	Panic in BL31 with register dump. System hung.	Capture the log, analyze the cause, and report to FAE if needed
INFO[UEFI]: eMMC init	eMMC driver is initialized	Informational and should always be printed
INFO[UEFI]: eMMC probed	eMMC card is initialized	Informational and should always be printed
ASSERT(UEFI): xxx : line-no	Runtime assert message in UEFI	Contact your NVIDIA FAE with this information. Usually the system is able to continue running.
INFO[UEFI]: PCIe enum start	PCIe enumeration start	Informational
INFO[UEFI]: PCIe enum end	PCIe enumeration end	Informational
ERR[UEFI]: Synchronous Exception at xxxxxx ERR[UEFI]: PC=xxxxxx ERR[UEFI]: PC=xxxxxx ...	UEFI Exception with PC value reported	Contact your NVIDIA FAE with this information
ERR[BL2]: FW auth failed	Image authentication error	Wrong image has been used in the current secure lifecycle. Switch to the correct image.
ERR[BL2]: IROT cert sig not found	Failed to load attestation certificates	Contact your NVIDIA FAE with this information

Message	Explanation	Action
ERR[BL2]: IROT cert sig not found	<p>Failed to load certification update record</p> <p>i Info Only relevant for certain BlueField-3 devices.</p>	Contact your NVIDIA FAE with this information
INFO[BL31]: PSC Turtle Mode detected	<p>PSC enters turtle mode</p> <p>i Info BlueField-3 only.</p>	Informational
INFO[BL31]: In Enhanced NIC mode	<p>BlueField-3 enters enhanced NIC mode</p>	Informational
ERR[BL31]: (set_page err pmbus_lsb err mfr_vr_mc err set_vout err)	<p>BlueField-3 power management programming error.</p> <p>i Info Usually happens when the I2C voltage regulator is not accessible.</p>	Contact your NVIDIA FAE with this information
INFO [BL31]: MB8: VDD adjustment complete	<p>BlueField-3 MainBin 8-core board VDD CPU adjustment</p>	Informational

Message	Explanation	Action
INFO [BL31]: VDD adjustment complete	BlueField-3 (non-8-core board) VDD CPU adjustment	Informational
INFO [BL31]: VDD: xxx mV	BlueField-3 VDD CPU voltage	Informational
ERR[BL31]: cannot access vr0 (or access vr1)	BlueField-3 unable to access voltage regulator (vr0 or vr1) via I2C	Contact your NVIDIA FAE with this information
ERR[BL31]: ATX power not detected!	ATX power is not connected	Contact your NVIDIA FAE with this information
INFO[BL31]: PTMERROR: Unknown OPN	Unable to detect the OPN on this device	Contact your NVIDIA FAE with this information
INFO[BL31]: PTMERROR: VR access error	Unable to access the voltage regulator on this device <div style="background-color: #ffffcc; padding: 10px; border: 1px solid #ccc;"> <p>i Info This also means power capping will be disabled.</p> </div>	Contact your NVIDIA FAE with this information
INFO[BL31]: power capping disabled	BlueField-3 power capping disabled	Informational
INFO[BL2]: boot mode (rshim emmc unknown)	Device boot mode (from external RShim or eMMC)	Informational
ERR[BL31]: ECC_SINGLE_ER ROR_CNT=xxx	Single ECC error counter report	Contact your NVIDIA FAE with this information
ERR[BL31]: ECC_DOUBLE_E RROR_CNT=xxx	Double ECC error counter report	Contact your NVIDIA FAE with this information

Message	Explanation	Action
ERR[BL31]: mss0 mss1: C0 C1 single-bit ecc, IRQ[%d]	MSS (0 or 1) channel (0 or 1) single-bit ECC error interrupt #	Contact your NVIDIA FAE with this information
ERR[BL31]: mss0 mss1: C0 C1 Double bit ecc, IRQ[%d]	MSS (0 or 1) channel (0 or 1) double-bit ECC error interrupt #	Contact your NVIDIA FAE with this information
ERR[BL31]: Double-bit ECC also detected in same buffer	Single/double ECC error detected in the same buffer	Contact your NVIDIA FAE with this information
ERR[BL31]: l3c: double-bit ecc	L3c double-bit ECC error detected	Contact your NVIDIA FAE with this information
ERR[BL31]: MSS%d DIMM%d single double bit ECC error detected	MSS DRAM single (or double) bit error detected	Contact your NVIDIA FAE with this information
ERR[BL31]: MSS%d SRAM double bit ECC error detected	MSS SRAM double bit ECC error detected	Contact your NVIDIA FAE with this information

IPMI Logging in UEFI

During UEFI boot, the BlueField sends IPMI SEL messages over IPMB to the BMC in order to track boot progress and report errors. The BMC must be in responder mode to receive the log messages.

SEL Record Format

The following table presents standard SEL records (record type = 0x02).

Byte(s)	Field	Description
1 2	Record ID	ID used to access SEL record. Filled in by the BMC. Is initialized to zero when coming from UEFI.

Byte(s)	Field	Description
3	Record Type	Record type
4 5 6 7	Timestamp	Time when event was logged. Filled in by BMC. Is initialized to zero when coming from UEFI.
8 9	Generator ID	This value is always 0x0001 when coming from UEFI
10	EvM Rev	Event message format revision which provides the version of the standard a record is using. This value is 0x04 for all records generated by UEFI.
11	Sensor Type	Sensor type code for sensor that generated the event
12	Sensor Number	Number of the sensor that generated the event. These numbers are arbitrarily chosen by the OEM.
13	Event Dir Event Type	[7] – 0b0 = Assertion, 0b1 = Deassertion [6:0] – Event type code
14	Event Data 1	[7:6] – Type of data in Event Data 2 <ul style="list-style-type: none"> • 0b00 = unspecified • 0b10 = OEM code • 0b11 = Standard sensor-specific event extension [5:4] – Type of data in Event Data 3 <ul style="list-style-type: none"> • 0b00 = unspecified • 0b10 = OEM code • 0b11 = Standard sensor-specific event extension [3:0] – Event Offset; offers more detailed event categories. See <i>IPMI 2.0 Specification</i> section 29.7 for more detail.
15	Event Data 2	Data attached to the event. 0xFF for unspecified. Under some circumstances, this may be used to specify more detailed event categories.

Byte(s)	Field	Description
16	Event Data 3	Data attached to the event. 0xFF for unspecified.

See *IPMI 2.0 Specification* section 32.1 for more detail.

Possible SEL Field Values

BlueField UEFI implements a subset of the IPMI 2.0 SEL standard. Each field may have the following values:

Field	Possible Values	Description of Values
Record Type	0x02	Standard SEL record. All events sent by UEFI are standard SEL records.
Event Dir	0b0	All events sent by UEFI are assertion events
Event Type	0x6F	Sensor-specific discrete events. Events with this type do not deviate from the standard.
Sensor Number	0x06	UEFI boot progress "sensor". If value is 0x06, the sensor type will always be "System Firmware Progress" (0x0F).

For Sensor Type, Event Offset, and Event Data 1-3 definitions, see next table.

Event Definitions

Events are defined by a combination of Record Type, Event Type, Sensor Type, Event Offset (occupies Event Data 1), and sometimes Event Data 2 (referred to as the Event Extension if it defines sub-events).

The following tables list all currently implemented IPMI events (with Record Type = 0x02, Event Type = 0x6F).

Note

Note that if an Event Data 2 or Event Data 3 value is not specified, it can be assumed to be Unspecified (0xFF).

Sensor Type	Sensor Type Code	Event Offset	Event Description, Actions to Take
System Firmware Progress	0x0F	0x00	System firmware error (POST error). Event Data 2: <ul style="list-style-type: none">0x06 – Unrecoverable EMMC error. Contact NVIDIA support.
		0x02	System firmware progress: Informational message, no actions needed. Event Data 2: <ul style="list-style-type: none">0x02 – Hard Disk Initialization. Logged when EMMC is initialized.0x04 – User Authentication. Logged when a user enters the correct UEFI password. This event is never logged if there is no UEFI password.0x07 – PCI Resource Configuration. Logged when PCI enumeration has started.0x0B – SMBus Initialization. This event is logged as soon as IPMB is configured in UEFI.0x13 – Starting OS Boot Process. Logged when Linux begins booting.

Reading IPMI SEL Log Messages

Log messages may be read from the BMC by issuing it a “Get SEL Entry Command” while it is in responder mode, either from a remote host, or from BlueField itself once it is

booted.

```
$ ipmitool sel list
7b | Pre-Init | 0000691604 | System Firmwares #0x06 | SMBus initialization | Asserted
7c | Pre-Init | 0000691604 | System Firmwares #0x06 | Hard-disk initialization | Asserted
7d | Pre-Init | 0000691654 | System Firmwares #0x06 | System boot initiated
$ ipmitool sel get 0x7d
SEL Record ID      : 007d
Record Type        : 02
Timestamp          : 01/09/1970 00:07:34
Generator ID       : 0001
EvM Revision       : 04
Sensor Type        : System Firmwares
Sensor Number      : 06
Event Type         : Sensor-specific Discrete
Event Direction    : Assertion Event
Event Data         : c213ff
Description        : System boot initiated
$ ipmitool sel clear
Clearing SEL. Please allow a few seconds to erase.
$ ipmitool sel list
SEL has no entries
```

ACPI BERT Logging

ACPI boot error record table (BERT) is supported to log last boot error in Linux. Once Linux `printk` is enabled (e.g., by adding "`kernel.printk=8`" to `/etc/sysctl.conf`), it will try to report the errors automatically for last boot. The following is an example of such error reports:

```
[ 2.635539] BERT: Error records from previous boot:
[ 2.640434] [Hardware Error]: event severity: fatal
[ 2.645331] [Hardware Error]: Error 0, type: fatal
[ 2.650236] [Hardware Error]: section type: unknown, c6adf9e6-1108-4760-8827-003d059fe2e1
[ 2.658606] [Hardware Error]: section length: 0x35
[ 2.663580] [Hardware Error]: 00000000: 52524520 4645555b 203a5d49 0a0d0a0d  ERR[UEFI]: ....
[ 2.672284] [Hardware Error]: 00000010: 636e7953 6e6f7268 2073756f 65637845  Synchronous Exce
[ 2.680987] [Hardware Error]: 00000020: 6f697470 7461206e 36783020 37313643  ption at 0x6C617
[ 2.689696] [Hardware Error]: 00000030: 34 37 30 0d 0a
```

...

© Copyright 2024, NVIDIA. PDF Generated on 08/20/2024