Release Notes

InfiniScale™ III MT47396 Firmware

fw-47396 Rev 1.0.5 -- Modifed on August 10, 2010

InfiniScale(TM) III 47396 Firmware Release Notes

Document Number:

Mellanox Technologies, Inc.
2900 Stender Way
Santa Clara, CA 95054
U.S.A.
www.mellanox.com

Tel: (408) 970-3400
Fax: (408) 970-3403

Mellanox Technologies Ltd
PO Box 586 Hermon Building
Yokneam 20692
Israel

Tel: +972-4-909-7200
Fax: +972-4-959-3245

1 Overview

These are the release notes for the MT47396 InfiniScale^{TM} III firmware fw-47396, Rev 1.0.5. This firmware complements the InfiniScale III MT47396 silicon architecture with a set of advanced features, allowing easy and remote management of the switch.

> Note: After burning new firmware to the InfiniScale III switch device, reboot the machine so that the new firmware can take effect.

This document includes the following sections:

1.1 Notes on This Release

- This firmware release complies with the *InfiniScale III MT47396 Programmer's Reference Manual (PRM), Doc. #2235PM, Rev 1.13,* except for features listed under Section 3, "Unsupported Features," on page 4.
- This firmware release supports the MT47396A1-FCC, MT47396A1-FDC, and MT47396A1-FDC-D InfiniScale III devices only. MT47396A1-FDC-D is the only device supporting IB ports operation at DDR (Double Data Rate).

1.2 Changes to This Document

- Corrected Table 10, "PLLr0/1Ldiv settings according to input reference clock frequency," on page 21. The order of the first two rows has been reversed, and the last row has been turned "reserved".

1.3 Requirements

- One of the following burning tools:
  - **spark** and **ibspark** applications (part of MFT ver. 2.0.0 or later)
  - **ibfwmgr** application (part of IBADM ver. 1.6.0 or later)
  - **is3burn** application (part of MST ver. 4.2.0 or later)
- Supported I2C cards:
  - MTUSB-1 (USB to I2C adapter)
  - ISA Calibre
- This firmware supports the MT47396 InfiniScale III switch device only

2 Changes and Major New Features

- New recovery flow for Detected and Uncorrectable Errors (DUE)
- Changed the ParityErrorFlow setting from HaltFW to SwReset in the INI files included in the FW package. See Section 7.2.4.12 on page 11 for the expected device behavior in each case.

3 Unsupported Features

- Baseboard Management Agent (BMA)

4 Bug Fixes

Table 1 - Bug Fixes

Index	Issue	Description
1.	SMA uses wrong PKey of received QP1 MAD when responding to a QP1 MAD	SMA now uses the corresponding PKey in its own table (and not the MAD's original PKey) for a received QP1 MAD when responding to a QP1 MAD
2.	MKey trap returns the actual MKey instead of Mkey=0	The MKey trap now returns Mkey=0
3.	Wrong 64-bit PortXmitData and PortRcv-Data counter values	When operating as 64-bit counters, and under stressful conditions, the PortXmitData and PortRcvData counters may incorrectly increase the high 32 counter bits by an extra 1
4.	OptimizedSLtoVLMappingProgramming IB specification 1.2.1 compliance	This field is now two bits wide (bits 95:94) in the SwitchInfo table as specified in Section: 14.2.5.4 SWITCHINFO of the *InfiniBand Architecture Specification, Release 1.2.1, Vo1. 1*. Previously, this field was one bit wide (bit 94).

5 Known Issues

Table 2 - Known Issues

Index	Issue	Description	Current Implemented Workaround in FW	Possible Workaround	Patch Release (fix)	Scheduled Release (fix)
1.	Cannot set PKey Table of the disabled sub-ports of a 3-port cluster configured as a single port	If a 3-port cluster is configured as a single 4X/12X port, then the PKey Table of two disabled sub-ports of this cluster cannot be changed by Set(PKeyTable) MAD. (ID: 16268, 20969)	NA	NA	NA	NA

6 SMA/GSA Attributes

The following tables summarize the attributes supported by the management agents provided in this release.

Table 3 - SMA Supported Attributes

Attribute	Support
Notice	X
NodeDescription	X
NodeInfo	X
SwitchInfo	X
GUIDInfo	X
PortInfo	X
Partition Key Table	X
SLtoVLMappingTable	X
VLArbitration	X
LinearForwardingTable	X
RandomForwardingTable	
MulticastForwardingTable	X
SMInfo	
VendorDiag	
LedInfo	

Table 4 - Performance Management Supported Attributes

Attribute	Support
ClassPortInfo	X
PortSamplesControl	X
PortSamplesResult	X
PortCounters	X
PortCountersExtended (for data counters only)	X

7 InfiniScale III Firmware Initialization And Configuration (.INI) File

The Mellanox firmware burning tools enable setting initialization and configuration attributes to suit a user's specific system by the use of a special (.INI) file. This section describes this user-supplied initialization and configuration file.

To begin with, the .INI file is a text file composed of several initialization and configuration *sections*. The user may choose to include all sections and all attribute settings in the final .INI file, and modify some of the attributes as required. Alternatively, the user may choose to keep only the sections with changes to the existing settings, with only those attributes that are to be modified.

This .INI file is described in the following sub-sections:

- ".INI File Format" on page 6
- "Description and Usage of .INI File Sections" on page 7

7.1 .INI File Format

The .INI file is actually a concatenation of (a part or all) section specific initialization and configuration settings. Each section in the .INI file starts with its name between square brackets, e.g. [EEPROM], [General], etc. The section name is followed by one or more lines of configuration settings and comments, as in the partial .INI file shown below. Note that comment lines start with a semicolon.

Example:

```
[EEPROM]
; This is a comment line
amount = 0x4
eeprom1_address = 0x56
eeprom1_size = 64

[General]
PortsBufferingMode = SAF
;This is another comment line
DEVID = 0xb924
VID = 0x2c9
AutoPowerSave = Enable

; End of (partial) .INI file
```

7.2 Description and Usage of .INI File Sections

The .INI file sections are:

7.2.1 [EEPROM]

The EEPROM section is used to indicate to the burning tool, and later to the firmware itself, the amount, size, slave address, and order of the EEPROMs on the board.

NOTE: The first boot EEPROM must be consistent with the number indicated through strapping pins, as described in the *InfiniScale III Hardware Reference Manual (HRM)*.

Example. The example below describes a board with 4 EEPROMs. The size of each EEPROM is 64KB. The I2C slave address of each EEPROM is defined as well.

```
[EEPROM]
;number of EEPROMs
```

amount = 0x4

;I2C slave address of EEPROM #1

eeprom1_address = 0x56

; size of EEPROM in KB

eeprom1_size = 64

eeprom2_address = 0x57

eeprom2_size = 64

eeprom3_address = 0x52

eeprom3_size = 64

eeprom4_address = 0x53

eeprom4_size = 64

7.2.2 [PS_INFO]

This section provides firmware configuration (ini) information used by the IB Administration tools. It has no impact on the generated firmware image.

Example:

Name = Reindeer4X

Version = 0.5.0

Description = Mellanox Reindeer switch with 4X ports

7.2.3 [PSID]

Parameter Set ID (PSID) describes a pre-defined set of parameters. It is used for quick identification of different board configurations. This attribute can be assigned a string built from up to 16 characters of the set: [A-Z,0-9].

Example:

PSID = MT_0060000001

7.2.4 [General]

The General section includes several general configuration attributes of the InfiniScale III device. These attributes are described in the following sub-sections.

7.2.4.1 INI file number

This attribute serves as a .INI file revision number. Must be a number of 32 bits.

Example:

INIFileNumber = 0

7.2.4.2 Buffering Mode

Set the attribute PortsBufferingMode to force *all IB ports of the device* to operate in one of two possible buffering modes: either Store and Forward (SAF) or Cut-Through (CT).

In case CT is chosen for buffering mode, then each IB port of the device will get dynamically configured so as to use the minimum buffer size possible. This size depends on the port link width.

In the [General] section write either

PortsBufferingMode = CT

or

PortsBufferingMode = SAF

7.2.4.3 Device and Vendor IDs

It is possible to configure the Device and Vendor IDs that will be used by the device as part of NodeInfo.

For InfiniScale III (MT47396) set Device ID to 0xb924 (= 4739610). In the [General] section write:

DEVID = 0xb924

For Mellanox Technologies set the Vendor ID to 0x2c9. In the [General] section write:

VID = 0x2c9

7.2.4.4 Power-Save

The SerDes devices can operate with Auto-Power-Save mode *enabled* or *disabled*. If enabled, the SerDes will be automatically powered-down if the physical link of the port is *disabled*. If this attribute is disabled, the SerDes will consume power whether the physical link of the port is enabled or not.

In the [General] section write either

AutoPowerSave = Enable

or

AutoPowerSave = Disable

7.2.4.5 GPIO-Led Enable

The GPIOLinkStatusLedEn attribute is used to enable or disable encoding of the link status of the ports on LEDs via GPIO pins of the device. For details, please refer to the *InfiniScale III (MT47396) Hardware Reference Manual*.

In the [General] section write either

GPIOLinkStatusLedEn = Enable

or

GPIOLinkStatusLedEn = Disable

7.2.4.6 SDO-Led Enable

The GMONLinkStatusGmonEn attribute is used to enable or disable encoding of the link status of the ports on LEDs via the serial System Monitoring pin (SDO) of the device. For details, please refer to the *InfiniScale III (MT47396) Hardware Reference Manual*.

In the [General] section write either

GMONLinkStatusGmonEn = Enable

or

GMONLinkStatusGmonEn = Disable

7.2.4.7 Led Blinking Enable

The LEDBlinkEn attribute is used to enable or disable encoding of the data transfer rate by *blinking* the LEDs. When disabled, the LEDs will only indicate whether the logical link is up or down.

In the [General] section write either

LEDBlinkEn = Enable

or

LEDBlinkEn = Disable

7.2.4.8 Single Led Indication Enable

The SingleLedIndication attribute is used to control the physical link status indication of the port cluster in Single mode. If enabled, a single led is used to indicate the physical link status (up or down). If disabled, in addition to the physical link status, two sub-port leds are used to indicate link width (1x, 4x, 12x).

In the [General] section write either

SingleLedIndication = Enable

or

SingleLedIndication = Disable

7.2.4.9 Trap Enable

The TrapFlagEn attribute is used to enable or disable the generation of all traps by the InfiniScale III device.

In the [General] section write either

TrapFlagEn = Enable

or

TrapFlagEn = Disable

7.2.4.10 GPIO Direction register

The InfiniScale III device has 64 GPIO pins. Each pin may be configured to act as input or output via the GPIO Direction register. If bit x of this GPIO Direction register is cleared (‘0), then GPIO pin number x will act as an input pin to InfiniScale III; if bit x is set (‘1), then GPIO pin number x will act as an output pin of InfiniScale III.

Example. To configure GPIO pins [31:16] as outputs, GPIO pins [63:32] and [15:0] as inputs, write in the [General] section:

GPIODir = 0xffff0000

Possible Conflict. If the attribute GPIOLinkStatusLedEn is set to *Enable*, the GPIO pins of InfiniScale III that are connected to LEDs should be set to act as outputs in the GPIO Direction register. Otherwise, there will be a conflict between the two configuration settings.

7.2.4.11 GPIO Polarity Register

The GPIO Polarity register is used to set the input polarity or output open-drain mode of the GPIO pins. Table 5 summarizes the possible configurations.

Table 5 - GPIO pin behavior based on Direction and Polarity register values

GPIO bit x Direction	GPIO bit x Polarity	GPIO pin x acts as
0	0	Non-Inverting Input
0	1	Inverting Input
1	0	Open Drain Non-Inverting Output
1	1	Open Drain Inverting Output

This feature is especially useful for applications where GPIO pins are used as interrupt inputs. Write in the [General] section the desired 64-bit value to GPIOPol.

Example. To set all GPIO pins as non-inverting inputs/outputs, write

GPIOPol = 0x0

7.2.4.12 Parity Error Flow

Describes the desired device behavior in case a Parity Error occurs in its internal memory. This attribute has two possible settings:

1. ParityErrorFlow = SwReset
   - This setting instructs the device to execute a software reset upon Parity Error.
2. ParityErrorFlow = HaltFW
   - Upon a Parity Error, this setting causes all IB links to go down, the GPIO pin defined by ParityErrorGpioBit (see below) to be asserted, and the firmware to halt.

7.2.4.13 Parity Error GPIO

Specifies the GPIO pin to be asserted in case of a firmware halt due to a Parity error. Possible values:

[0..15] if CpuMode=CPU

[0..15, 32..63] if CpuMode=GPIO

See explanation about CpuMode in “CPU Pins Mode” on page 14.

Example:

ParityErrorGpioBit = 0x0

7.2.4.14 Switch Health Status GPIOs

GPIO9 & GPIO10 can be used to report switch health status by enabling the paramter GpioStatusEnable.

In the [General] section write either

GpioStatusEnable = Enable

or

GpioStatusEnable = Disable

Note: If GpioStatusEnable is enabled, GPIO9 and GPIO10 will be configured as outputs.

7.2.4.15 SwitchRelayErrorMsk for Port 0 PortCauseBits MAD Entries

For **Port 0 only**, it is possible to mask the discarding of received packets that could not be forwarded by the switch relay. This is achieved using the *SwitchRelayErrorMsk* parameter. If *SwitchRelayErrorMsk* is enabled, then the value returned in the field *PortRcvSwitchRelayErrorCounter* in a PM (perfrormance) MAD response will always be zero.

In the [General] section write either

SwitchRelayErrorMsk = Enable

or

SwitchRelayErrorMsk = Disable

7.2.5 [Special]

The special section is used for configuring items which have special implications on firmware.

7.2.5.1 System Image GUID

System image GUID is a 64-bit number. Use sysimage_GUID attribute to set it.

Example: To set the system image GUID to 0, write in the [Special] section

sysimage_GUID = 0x0

7.2.5.2 Node Description

Node description is a string written to the node_descr attribute of the Subnet Management class.

Example: Describe the InfiniScale III device as follows in the [Special] section:

node_descr = MT47396 InfiniScale-III Mellanox Technologies

7.2.5.3 Node GUID

Node GUID is a 64-bit number used for NodeInfo. Set it using the NodeGUID attribute.

Example:

Node_GUID = 0x2c90000000000

7.2.5.4 Board Serial Number

The Board Serial Number (BSN) attribute can be used only if the board is supplied by Mellanox Technologies. (Please consult your FAE to know which Mellanox boards are currently supported.)

For example, if the board is an MTS2400, then this attribute should indicate the unique serial number of the board about to burnt. In the case the attribute is not applicable to your board, the following line should be left unchanged in the .INI file:

BSN = NO_BSN

7.2.5.5 Vendor Specific Enable

The attribute VendorSpecEn is used to enable the InfiniScale III Firmware to respond to Vendor Specific Class attributes, or disable it from responding. It determines the value of the bit IsVendorClassSupported in the Capability Mask register.

Write in the [Special] section either

VendorSpecEn = Enable

or

VendorSpecEn = Enable

7.2.6 [ENP0]

This section is dedicated for Enhanced Port 0 related attributes.

7.2.6.1 Enhanced Switch Port 0 Enable

If ENP0En attribute is set to Enable, the management port of the InfiniScale III device will function as Enhanced Port 0. If set to Disable, the management port will function as Base Port 0. This attribute also determines the value of the SwitchInfo.ENP0 bit.

Write in the [ENP0] section either

ENP0En = Enabled

or

ENP0En = Disabled

7.2.6.2 CPU Master Target Port Size

The attribute CPUMPortSize is used to indicate to the InfiniScale III CPU bus Master what is the port size of the target on the CPU bus.

For a 32-bit port size write

CPUMPortSize = T32

For a 16-bit port size write

CPUMPortSize = T16

7.2.6.3 CPU Interrupt Request

The attribute CPUIrq is used to indicate to the host which of its interrupt requests is connected to the InfiniScale III interrupt output.

Possible values are 0x0 - 0x5.

Example: To map the InfiniScale III interrupt output to host interrupt request no. 2 write:

CPUIrq = 0x2

7.2.7 [MISC]

This section combines miscellaneous attributes.

7.2.7.1 I2C Primary and Secondary slave addresses and enables

Use the I2CSlavePrimaryAddr attribute to set the slave address of the InfiniScale III device on the primary I2C bus; and use I2CSlavePrimaryEnable to enable this slave.

Use the I2CSlaveSecondaryAddr attribute to set the slave address of the InfiniScale III device on the secondary I2C bus; and use I2CSlaveSecondaryEnable to enable this slave.

Note: Setting the primary slave address overrides the information from the strapping pins.

Possible values for the I2C slave enables are: 0x0 or 0x1.

Examples:

I2CSlavePrimaryAddr = 0x6c

I2CSlavePrimaryEnable = 0x1

I2CSlaveSecondaryAddr = 0x2a

I2CSlaveSecondaryEnable = 0x1

7.2.7.2 CPU Pins Mode

The InfiniScale III pins that are used for PPC bus interface can also be used as GPIO pins. The CpuMode attribute allows the user to configure these pins as CPU or GPIO pins in advance.

Write in the [MISC] section either

CpuMode = Cpu

or

CpuMode = Gpio

7.2.7.3 System Monitoring Chain Disable

The SystemMonitoringDis attribute is used to disable the System Monitoring serial chain (via SDO pin). If disabled, the SDO output pin will remain undriven by the device, i.e. in the High-Z state.

Write in the [MISC] section either

SystemMonitoringDis = false

or

SystemMonitoringDis = false

7.2.7.4 System Monitoring Clock Cycle Time

The SystemMonitoringClCyT attribute is used to determine the System Monitoring chain clock cycle time. The Core clock cycle time is divided by the value specified by this attribute +1.

Range of values: 10 through 255

Example: If the following line is written in the [MISC] section, then the cycle time of the System Monitoring chain will be the core clock cycle time divided by 0x13 + 1 = 2010.

SystemMonitoringClCyT = 0x13

That is, System Monitoring Cyc Time = (Core clock cycle time) / 20.

7.2.7.5 I2C Bus Master Frequency (Primary and Secondary)

The PriI2CBusFreq and SecI2CBusFreq attributes are used to determine the bus frequency of the InfiniScale III device as a master on the primary and secondary I2C busses respectively. The number specified is in KHz units.

Range of values: 50 through 300.

Examples:

PriI2CBusFreq = 300

SecI2CBusFreq = 0x41

7.2.7.6 CPU Bus Request Mode

This attribute sets the ***CBR#*** pin output mode: regular active

For regular active output mode (1 for inactive, 0 for active) write:

CPUBusReqMode = 0x0

For Open Drain output mode write:

CPUBusReqMode = 0x1

Default output mode: Open Drain

This attribute is applicable only for A1 silicon devices. In A0 devices the ***CBR#*** pin is an Open Drain output, regardless of the value of this attribute.

7.2.7.7 Target Acknowledge Timer

Specifies the period of time, measured in cycles of CPU bus clock, the CPU bus master waits for a Target Acknowledge signal (CTA#, CTEA#, CRETRY#) after it asserted the Transaction Start signal (CTS#). When this timer expires the transaction is terminated with error status 4. If this attribute is set to 0 the timer is disabled.

Possible values: [0..0xFFF].

Example:

TATimer = 0x0

7.2.8 [IB_TO_HW_MAP]

This section describes the connectivity between the InfiniBand **silk/front-panel** (Logical IB) ports and the InfiniScale III switch device **hardware** (HW) ports. The connectivity, however, is described indirectly: instead of mapping each IB port to the InfiniScale III HW port it physically connects to, each IB port is mapped to an InfiniScale III **Internal** port. Each such Internal port is associated with a single HW port. As an example, see Figure 1 which depicts the mapping of Logical IB Port 1 to InfiniScale III HW Port 18 via Internal Port 15.

Figure 1: Mapping a Chassis Logical IB Port to a HW Port Using an Internal Port

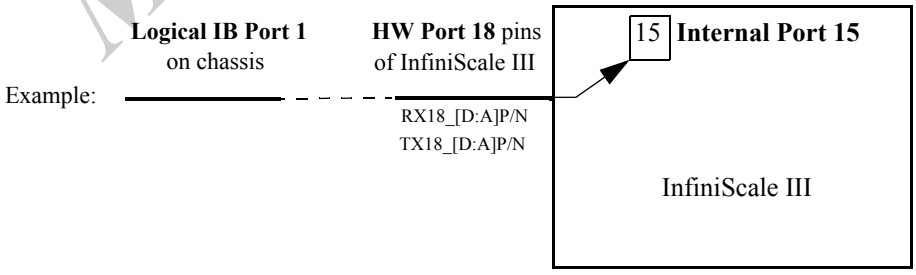

The mapping described in the figure will be set using the following set command in the IB_TO_HW_MAP section:

PORT1 = 15

where PORT1 indicates Logical IB Port 1 on the chassis, and 15 is Internal port associated with HW Port 18 of the InfiniScale III switch device. The [IB_TO_HW_MAP] section should contain 24 mappings: PORT1 - PORT24.

Refer to Table 6 to identify for each InfiniScale III **HW** port the **Internal** port it is associated with.

Table 6 - InfiniScale III HW Ports to Internal Ports Map

InfiniScale III HW Port	InfiniScale III Internal Port	InfiniScale III HW Port	InfiniScale III Internal Port
1	22	2	23
3	24	4	1
5	2	6	3
7	4	8	5
9	6	10	7
11	8	12	9
13	10	14	11
15	12	16	13
17	14	18	15
19	16	20	17
21	18	22	19
23	20	24	21

7.2.9 [IB_TO_LED_MAP]

For each Logical IB port on the chassis corresponds a (green or yellow) LED. The InfiniScale III employs 6 GPIO pairs (GPIO 16-27) and four GPIO qualifier signals (GPIO 28-31) to light up to 24 chassis IB port LEDs. This section defines the mapping between the Logical (chassis) IB port number and its GPIO control pins.

The mapping scheme to be presented here allows different board designs to have various connectivity and control schemes in a very simple and flexible manner.

Table 7 below provides code numbers by which it is possible to map logical port to their corresponding LED GPIO controls. Each code in a table cell defines a pair of GPIO pins and a qualifying GPIO pin.

Table 7 - Code numbers for mapping Logical (chassis) IB ports to their corresponding LED GPIO control pins

GPIO Pair \ Qualifier	GPIO28	GPIO29	GPIO30	GPIO31
GPIO 16,17	1	7	13	19
GPIO 18,19	2	8	14	20
GPIO 20,21	3	9	15	21
GPIO 22,23	4	10	16	22
GPIO 24,25	5	11	17	23
GPIO 26,27	6	12	18	24

Example 1: If Logical IB port 1 is controlled by the pair [GPIO26, GPIO27] and the qualifier GPIO30, then the [IB_TO_LED_MAP] section should include the following line:

PORT1 = 18

Example 2: If Logical IB port 24 is controlled by the pair [GPIO22, GPIO23] and the qualifier GPIO28, then the [IB_TO_LED_MAP] section should include the following line:

PORT24 = 4

7.2.10 [PortDisable]

This section is used to enable or disable the IB ports of the InfiniScale III, or to activate electronic keying (EKEYING) of IB ports.

7.2.10.1 Port Disable/Enable

The PortDisable section allows enabling or disabling specific ports (in the range 1-24). If a port is disabled here, then its physical link will remain disabled after boot, and its SerDes TX will be powered down.

Note: By default, all InfiniScale III ports are enabled in the .INI file.

In the [PortDisable] section, write:

PORT*i* = Disabled to disable PORT*i, i=1..24*.

PORT*i* = Enabled to enable PORT*i, i=1..24*.

7.2.10.2 EKEYING - Electronic Keying

E-keying an IB port causes this port to indicate it is in 'Polling' state. In reality the port is disabled, however reporting a 'Polling' state prevents moving the port to 'Active' state.

Using the .INI file, it is possible to set the e-keying state of InfiniScale III ports either statically or dynamically. In the first case, the IBEKEYing parameter is programmed with the desired e-keying state of each port upon exiting boot. See "E-keying a Port Upon Exit from Boot:" on page 17.

In the second (dynamic) case, the e-keying state of InfiniScale III ports is set by reading from hardware e-key registers on the board. To trigger a read to these registers, however, an interrupt must occur on GPIO15. See "Defining E-keying Registers for Access upon a GPIO15 Interrupt:" on page 17.

E-keying a Port Upon Exit from Boot:

The parameter IBEKEYing is 24 bits wide, with each bit represeting a port number in the range 1..24. Bit 0 represents IB port 1 of the InfiniScale III, and bit 23 represents IB port 24. If port *i* is desired to wake in 'Polling' state upon exiting boot, set bit *i-1* in this parameter to 1. Thus setting IBEKEYing to 0xB6 in the .INI file marks ports 2, 3, 5, 6 and 8 for ekeying.

Note: By default no port is set for ekeying, i.e. IBEKEYing =0.

Defining E-keying Registers for Access upon a GPIO15 Interrupt:

The IB ports of the InfiniScale III are divided into three ekeying groups: 1-8, 9-16 and 17-24. For each of these goups three parameters are defined as summarized in Table 8.

Table 8 - E-key Register Access Parameters

Parameter Type	Parameter Name	Description
E-key Register Access Enable	EKEYingPorts1_8_en	Enables / Disables access[1] to the E-key Register for ports 1-8
	EKEYingPorts9_16_en	Enables / Disables access[1] to the E-key Register for ports 9-16
	EKEYingPorts17_24_en	Enables / Disables access[1] to the E-key Register for ports 17-24
I^2C Address of E-key Register	EKEYingPorts1_8_devid	I^2C address of E-key Register for ports 1-8
	EKEYingPorts9_16_devid	I^2C address of E-key Register for ports 9-16
	EKEYingPorts17_24_devid	I^2C address of E-key Register for ports 17-24
Offset within E-key Register	EKEYingPorts1_8_address	Byte offset within E-key Register for ports 1-8 (offset from EKEYingPorts1_8_devid)
	EKEYingPorts9_16_address	Byte offset within E-key Register for ports 9-16 (offset from EKEYingPorts9_16_devid)
	EKEYingPorts17_24_address	Byte offset within E-key Register for ports 17-24 (offset from EKEYingPorts17_24_devid)

1. If an E-key Register does not exist, this parameter should be set to Disabled.

By default, all 'E-key Register Access Enable' parameters are set to Disabled (i.e., no E-key Registers are accessible upon exiting boot).

Upon an interrupt (rising edge) on GPIO15, the enabled E-key Registers are read.

Example:

Following is an excerpt from a possible [PortDisable] section in a .ini file where:

- Ports 1-8 have no E-key register
- Ports 9-16 and 17-24 share the E-key device (same I2C address) but the register offsets are different
- Upon an interrupt on GPIO15 (rising edge), the E-key Register for ports 9-16 and the E-key Register for ports 17-24 will be read.

```
;Disable access to E-key register for ports 1-8
EKEYingPorts1_8_en = Disabled
;I2C Address (device id) of E-key register for ports 1-8
EKEYingPorts1_8_devid = 0x0
;Offset of ekeying register for ports 1-8
EKEYingPorts1_8_address = 0x0

;Enable access to E-key register for ports 9-16
EKEYingPorts9_16_en = Enabled
;I2C Address (device id) of E-key register for ports 9-16
EKEYingPorts9_16_devid = 0x21
;Offset of ekeying register for ports 9-16
EKEYingPorts9_16_address = 0x0
```

;Enable access to E-key register for ports 17-24

EKEYingPorts17_24_en = Enabled

;I2C Address (device id) of E-key register for ports 17-24

EKEYingPorts17_24_devid = 0x21

;Offset of ekeying register for ports 17-24

EKEYingPorts17_24_address = 0x1

7.2.11 [LinkWidthSupp]

The 24 IB ports of the InfiniScale III are viewed as eight 3-port clusters. Each such cluster gets a single command line in the [LinkWidthSupp] section to configure the link widths of its 3 ports. The configuration settings in this section determine both the LinkWidthSupported and the LinkWidthEnabled attributes of all InfiniScale III ports during boot. See Section 7.2.11.1 for more details.

The clusters are numbered CL1 through CL8, and the ports are assigned to these clusters in an ascending order. Thus, cluster CL1 groups ports 1, 2, and 3; cluster CL2 groups ports 4, 5, and 6;..., cluster CL8 groups ports 22, 23, and 24.

To configure a cluster with 3 ports, write in the [LinkWidthSupp] section lines in the following format:

CL*i* = config_value

where *i* represents a cluster number in the range 1,2,...,8, and config_value represents one of the values described in the table below.

Table 9 - Configuration Values for Cluster Port Link Widths

config_value	Resulting Cluster Configuration
TRIO_1X	the cluster is configured as 3 separate ports, each limited to 1X link width only
TRIO_1X4X	the cluster is configured as 3 separate ports, each supporting 1X or 4X link width
TRIO_4X	the cluster is configured as 3 separate ports, limited to 4X link width only
SINGLE_1X	the cluster is configured as a single port, limited to 1X link width only
SINGLE_1X4X	the cluster is configured as a single port supporting 1X and 4X link width
SINGLE_1X4X12X	the cluster is configured as a single port supporting 1X, 4X and 12X link widths
AUTOCONF_1X4X12X	the cluster will auto-configure itself to match the configuration of the (peer) cluster connected to it through the IB cable

Note that in case a cluster is set as AUTOCONF_1X4X12X, then it may rise as a single 12X port or as 3 4X ports, depending on the peer cluster capabilities at the time of link bring-up. This may occur in every instance of bring-up. Thus the cluster may start out as 3 4X ports prior to link (re)establishment, and turn into a single 12X port after (re)establishment, or vice versa.

Examples:

CL1 = TRIO_1X4X

CL2 = TRIO_1X4X

CL3 = SINGLE_1X4X

CL4 = SINGLE_1X

CL5 = TRIO_1X4X

CL6 = SINGLE_1X4X12X

CL7 = TRIO_1X4X

CL8 = AUTOCONF_1X4X12X

7.2.11.1 Link Width Lockup

An IB port in a trio configured to 4X width (4-SerDes port) by means of the CL*i* parameter above may still be switched by the Subnet Manager to 1X width (1-SerDes port). To prevent this scenario, the user can enable the configuration parameter *LinkWidthLockup*. This indicates to the SM that a device port configured as a 4X port will not reconfigure to a 1X port even if asked to by the SM.

Possible values for the *LinkWidthLockup* parameter are Disable (the default) and Enable.

> Note: This mode of operation, where LinkWidthLockup is enabled, serves as an extension to the functionlities defined in the *InfiniBand Architecture Specification (Release 1.2)*.

Example: Under the following parameters setting, the three ports in trio 2 will be configured to 4X width each, and will *not* reconfigure to 1X upon a request by the SM.

LinkWidthLockup = Enable
CL2 = TRIO_4X

7.2.12 [PLL]

This section is used to set the PLL boot record values.

7.2.12.1 Serial EEPROM Clock Rate Divider

When a serial EEPROM is present, the attribute ClockRateDivider is used to set the desired serial clock rate for reading the remaining part of the PLL Configuration Record via the I2C bus. It is an 8-bit number that the InfiniScale III multiplies by 16 to obtain the number by which to divide the external clock frequency. Thus, the I2C serial clock frequency will be given by the following equation:

Serial_Clock_Frequency = External_Clock_Frequency / (16 x ClockRateDivider)

In the [PLL] section, write a line in the following format:

ClockRateDivider = number

where number is an integer between 1 and 255.

7.2.12.2 Core PLL Divider

The core clock PLL in the InfiniScale III includes a divider for the external clock input. This divider requires two parameters, M and N, to produce required output clock frequency. The parameter M is supplied by the ClockDividerNumerator attribute, and N is supplied by the ClockDividerDenominator attribute in the [PLL] section of the INI file.

The core PLL output clock frequency is calculated by the following equation:

Fout = [M/(2 x N)] x Fin

where,

Fout is the core PLL output clock frequency

Fin is the external input clock frequency

N is a 3-bit wide Clock Divider Denominator

M is a 4-bit wide Clock Divider Numerator

For example, if the external clock input to InfiniScale III runs at 66MHz, i.e. Fin = 66, and the divider parameters are N=1 and M=5, then the PLL output core clock frequency will be:

Fout = [5/(2 x 1)] x 66 = 165MHz.

In the [PLL] section the attribute settings will appear as follows (for the example above):

ClockDividerDenominator = 0x1

ClockDividerNumerator = 0x5

7.2.12.3 VCO Range

The VCO_RANGE attribute should be to 1 to indicate that the VCO frequency is greater than or equal to 133MHz. Otherwise, it should be set to 0.

Example:

VCO_RANGE = 0x1

7.2.12.4 SerDes PLL Ratio – LDIV

InfiniScale III has two PLLs which provide the high speed differential clocks to the SerDes devices. PLLr0 provides the clock to ports 1 through 12, while PLLr1 provides the clock to ports 13 through 24.

The PLLr0Ldiv and PLLr1Ldiv attributes are used to set the ratio between the frequency of the external refclk pins V0RCLKP/N and V1RCLKP/N and the SerDes differential clock outputs. The SerDes differential clock nominal frequency is 2.5GHz, whether the SerDes operates at 2.5Gb/s (SDR) or 5Gb/s (DDR). For setting IB port operation at DDR, see parameters ending with the string “Speed5G” and the *IB DDR Auto-Negotiation Application Note, Document no. 2162AN*.

In order to obtain the nominal 2.5 GHz clock, the frequency of the external refclk is summarized in the following table:

Table 10 - PLLr0/1Ldiv settings according to input reference clock frequency

PLLr0/1Ldiv	Required External Reference Clock Frequency
000	125MHz
001	250MHz
010	62.5MHz
011-111	Reserved

Example: With a 250MHz reference input clock, set the attributes in the [PLL] section to 001 to obtain 2.5 GHz link speed.

PLLr0Ldiv = 0x1

PLLr1Ldiv = 0x1

7.2.12.5 Monitoring of the SerDes PLL Reference Clock

SerDes PLL Reference Clock can be monitored on the special device pin VPTEST. There are two parameters defining which of the PLL's clocks is issued on this pin: Er0Out and Er1Out, as summarized in the following table:

Table 11 - Selection of RefClk for Monitoring

Er0Out	Er1Out	Reference clock driven out
0	0	VPTEST pin is Hi-Z
0	1	RefClk1
1	0	RefClk0
1	1	VPTEST pin is Hi-Z

Example:

Er0Out = 0x0

Er1Out = 0x0

7.2.12.6 Core PLL Stabilization Time

The PLLStabilizationTime attribute is used to define the core PLL stabilization time. The value supplied indicates the number of external clock cycles required for PLL stabilization. During this stabilization period, the internal reset signal of InfiniScale III will remain asserted.

In the [PLL] section write a line in the following format:

PLLStabilizationTime = number

where number is an integer *larger* than 0.

7.2.13 [REV_LANE]

Each port of the InfiniScale III has 4 lanes for data transfer. This section is used to configure the Lane Reversal feature of these IB ports, taking into account the 3-port clustering mentioned in Section 7.2.11, "[LinkWidthSupp]," on page 19. Due to differences between receive (RX) and transmit (TX) attributes, they are discussed here separately.

7.2.13.1 RX attributes

The RX lines of each IB port have two attributes for setting in the [REV_LANE] section: Rx_Rev_lane_IB_port*X* and Rx_Sub_Rev_lane_IB_port*X*, where *X* is actually a cluster number [1..24].

The Rx_Rev_lane_IB_port*X* attribute. The usage of this attribute depends on the cluster configuration:

3. If a port is part of a cluster is in "Trio" mode, i.e. it consists of 3 separate ports, each supporting 1X or 4X link widths, then the Rx_Rev_lane_IB_port*X* attribute affects only the referenced port, and determines whether the 4 lanes of the port are to be reversed, not reversed, or automatically configured by firmware. To configure such a port in the cluster using this attribute, write in the [REV_LANE] section a line in the following format:

   Rx_Rev_lane_IB_port*X* = AutoConfig | Off | On

   where one of the 3 optional values is chosen. If AutoConfig is chosen, the firmware will configure the lanes; Off will keep the order 0..3; On will reverse it to 3..0.

   Note: It is highly recommended to set this RX parameter in this section to "AutoConfig".

4. If a port is part of a cluster in “Single” mode, i.e. it is a single IB port supporting up to 12X link width, then the Rx_Rev_lane_port*X* attribute determines whether *three ports* in the cluster are to be reversed, not reversed, or automatically configured by firmware. The term “reversed” in this context refers to port-ordering but not lane-ordering. This attribute has no effect on the lanes within each of these three ports. Thus, if the value “On” is set to Rx_Rev_lane_port*X* attribute, and the cluster ports are p1, p2, and p3, their order will be reversed to p3, p2, and p1. To configure the lanes of each port, see the next RX attribute.

The Rx_Sub_Rev_lane_IB_port*X* attribute. The second RX parameter in the [REV_LANE] section is used only when a cluster is in “Single” mode; it has no effect on a cluster in “Trio” mode. When in “Single” mode, this parameter determines whether the lanes of a specific port of the 3-port 12X cluster are to be reversed, not reversed, or automatically configured by firmware. If the Rx_Sub_Rev_lane_IB_port*X* attribute is “On”, then the lanes 0..3 of the port will be reversed to 3..0.

Note: It is highly recommended to set this RX parameter in this section to “AutoConfig”.

Example: Suppose we want to configure the RX attributes of a cluster comprised of ports 7, 8, and 9, and set it in SINGLE_1X4X12X mode. To have the cluster ports reversed in order and the port-lanes reversed in order, we would set the RX attributes as follows:

Rx_Rev_lane_IB_port*7* = On

Rx_Sub_Rev_lane_IB_port7 = On

Rx_Sub_Rev_lane_IB_port8 = On

Rx_Sub_Rev_lane_IB_port9 = On

Note that when setting the Rx_Rev_lane_IB_port*X* attribute for a single-mode cluster, *X* must be the Master Port of the cluster (i.e. the smallest port number in the cluster, 7 in this example). However, it is necessary to specify for each port separately whether its lanes should be reversed or not. Thus, to keep the lanes of port 8 not reversed, we should set the attributes as follows:

Rx_Rev_lane_IB_port*7* = On

Rx_Sub_Rev_lane_IB_port7 = On

Rx_Sub_Rev_lane_IB_port8 = Off

Rx_Sub_Rev_lane_IB_port9 = On

7.2.13.2 TX attribute

Tx_Rev_lane_port*X*. This attribute is used in a similar manner to the Rx_Rev_lane_port*X* (See above). The reversal of port lanes of a cluster in “Single” mode is automatically handled by firmware in accordance with the configuration of Section 7.2.8, “[IB_TO_HW_MAP],” on page 15.

To set this attribute in the [REV_LANE] section write a line in the following format:

Tx_Rev_lane_IB_port*X* = On | Off | AutoConfig

Tx_Trio_Rev_lane_IB_cluster*Y*. This attribute is applicable only in A1 silicon devices. It allows TX lane reversal in a cluster configured in TRIO mode in the manner described in the following table:

Table 12 - TX Trio Lane Reversal

Original lanes	Reversed lanes
0..3	11..8
4..7	4..7 (no change)
8..11	3..0

Possible values: Off, On

Example:

Tx_Trio_Rev_lane_IB_cluster*Y* = On

7.2.14 [Polarity]

This section is used to configure the polarity of the RX and TX lines of the SerDes interfaces for each IB port of the InfiniScale III device. Due to differences between RX and TX attributes, they are discussed here separately.

7.2.14.1 RX Polarity

The polarity of RX port lines can either be forced in the .INI file or left to be auto-configured. There are two attributes in this regard: Rx_Force_Pol_IB_port*X* and Rx_Pol_IB_port*X*, where *X*=1, 2,...24. The attribute Rx_Force_Pol_IB_port*X* determines whether the RX polarity is be forced in the INI file or to be auto-configured: if it is to be forced, then it requires setting the second attribute Rx_Pol_IB_port*X* as well; otherwise, Rx_Pol_IB_port*X* will be ignored.

Rx_Force_Pol_IB_port*X* can assume one of two values: Enable or Disable. To force a polarity, set this attribute to Enable; to have it auto-configured, set the attribute to Disable.

Rx_Pol_IB_port*X* can assume one of 16 possible values: 0x0 through 0xf. Each of the four bits determines the polarity of one RX lane of the port. A ‘0 bit for a lane will have it operate in normal (active high) polarity, whereas a ‘1 bit will have it operate in reverse (active low) polarity.

Examples:

1. To have the polarity of RX lanes of port 5 *auto-configured*, write in the [Polarity] section the following line:

   Rx_Force_Pol_IB_port5 = Disable

2. To *force* the polarity of all four RX lanes of port 5 to act in reverse polarity, write in the [Polarity] section the following lines:

   Rx_Force_Pol_IB_port5 = Enable

   Rx_Pol_IB_port5 = 0xf

3. To *force* the polarity of RX lanes 0, 1, and 3 of port 5 to act in normal polarity, and lane 2 in reverse polarity, write in the [Polarity] section the following lines:

   Rx_Force_Pol_IB_port5 = Enable

   Rx_Pol_IB_port5 = 0x4

7.2.14.2 TX Polarity

The polarity of TX port lines is always forced in the INI file. For each port there 4 lanes, and each lane cab be set to operate in normal (active high) polarity, or in reverse (active low) polarity. In this case, one attribute is required per port, i.e. Tx_Pol_IB_Port*X*, where *X*=1, 2,...,24. This attribute takes the same range of values as in the case of its RX parallel.

Example: To *force* the polarity of TX lanes 0 and 1 of port 9 to act in normal polarity, and lanes 2 and 3 in reverse polarity, write in the [Polarity] section the following line:

Rx_Pol_IB_port9 = 0xc

7.2.15 [SERDES_Equa_CFG]

This section is optionally used to configure the SerDes Equalization level in Single Data Rate (SDR) mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume an equalization level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x0.

Example: To set lane 2 of port 3 to level 7, write in the [SERDES_Equa_CFG] section the following line:

PORT3_SER2 = 0x7

7.2.16 [SERDES_Equa_CFG_Speed5G]

This section is optionally used to configure the SerDes Equalization level in Double Data Rate (DDR) mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of the InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume an equalization level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x0.

Example: To set lane 2 of port 3 to level 7, write in the [SERDES_Equa_CFG_Speed5G] section the following line:

PORT3_SER2 = 0x7

7.2.17 [SERDES_OutPut_Voltage]

This section is optionally used to tune the SerDes output voltage level in SDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume an output voltage level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x2.

Example: To set lane 1 of port 9 to output voltage level 0xb, write in the [SERDES_OutPut_Voltage] section the following line:

PORT9_SER1 = 0xb

7.2.18 [SERDES_OutPut_Voltage_Speed5G]

This section is optionally used to tune the SerDes output voltage level in DDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume an output voltage level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x2.

Example: To set lane 1 of port 9 to output voltage level 0xb, write in the [SERDES_OutPut_Voltage_Speed5G] section the following line:

PORT9_SER1 = 0xb

7.2.19 [SERDES_Pre_Amp_OutPut]

This section is optionally used to tune the SerDes Pre-Amplifier output voltage level in SDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume a pre-amplifier output voltage level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x2.

Example: To set lane 0 of port 23 to pre-amplifier output voltage level 0xe, write in the [SERDES_Pre_Amp_OutPut] section the following line:

PORT23_SER0 = 0xe

7.2.20 [SERDES_Pre_Amp_OutPut_Speed5G]

This section is optionally used to tune the SerDes Pre-Amplifier output voltage level in DDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume a pre-amplifier output voltage level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x2.

Example: To set lane 0 of port 23 to pre-amplifier output voltage level 0xe, write in the [SERDES_Pre_Amp_OutPut_Speed5G] section the following line:

PORT23_SER0 = 0xe

7.2.21 [SERDES_Pre_Emp_Out]

This section is optionally used to tune the SerDes Pre-Emphasis output level in SDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume a pre-emphasis output level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x0.

Example: To set lane 3 of port 19 to pre-emphasis output level 0x6, write in the [SERDES_Pre_Emp_Out] section the following line:

PORT19_SER3 = 0x6

7.2.22 [SERDES_Pre_Emp_Out_Speed5G]

This section is optionally used to tune the SerDes Pre-Emphasis output level in DDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume a pre-emphasis output level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x0.

Example: To set lane 3 of port 19 to pre-emphasis output level 0x6, write in the [SERDES_Pre_Emp_Out_Speed5G] section the following line:

PORT19_SER3 = 0x6

7.2.23 [SERDES_Pre_Emp_Pre_Amp]

This section is optionally used to tune the SerDes Pre-Emphasis level of the Pre-Amplifier in SDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume a pre-emphasis level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x0.

Example: To set lane 2 of port 15 to pre-emphasis level 0x4 of pre-amplifier, write in the [SERDES_Pre_Emp_Pre_Amp] section the following line:

PORT15_SER2 = 0x4

7.2.24 [SERDES_Pre_Emp_Pre_Amp_Speed5G]

This section is optionally used to tune the SerDes Pre-Emphasis level of the Pre-Amplifier in DDR mode. Each of the 96 SerDes devices (24 ports x 4 lanes) of InfiniScale III is configured separately. The attribute used for each lane is PORTi_SERj, where i represents a port number (1,2,...24), and j represents a lane number (0,1,2,3).

PORTi_SERj can assume a pre-emphasis level number between 0x0 and 0xf. By default, if this attribute is not specified for some lane(s) in the .INI file, it will hold the value 0x0.

Example: To set lane 2 of port 15 to pre-emphasis level 0x4 of pre-amplifier, write in the [SERDES_Pre_Emp_Pre_Amp_Speed5G] section the following line:

PORT15_SER2 = 0x4

7.2.25 [Credits_Time]

This section is optionally used to set, for each InfiniScale III IB port, the maximum and minimum elapsed time periods between two successive transmissions of credit packets on each Virtual Lane (VL). The two attributes used are MAX_TIME_PORT*i* and MIN_TIME_PORT*i*, where *i* is a port number (1,2,...24).

MIN_TIME_PORT*i* is specified in units of (64 x symbol-time), where symbol-time equals 256ns. After this minimum time has elapsed, the output port *i* schedules for transmission a credit-update packet for every VL. The transmission occurs once new credits become available.

Note: In case of a DDR IB port, the MIN_TIME_PORT*i* value will be divided by 2.

MAX_TIME_PORT*i* is used to force the delivery of credit packets for every VL after a certain amount of time has elapsed (even if no new credits are available). MAX_TIME_PORT*i* is specified in units of "cycles of the minimum time counter". This means that the actual maximum elapsed time between successive credit packets will be given my the following expression: PORT_MAX_TIME*i* x PORT_MIN_TIME*i* x (64 x symbol time).

Note that in the case where MAX_TIME_PORT*i* = 0, no VL credit packets will be transmitted as long as *no new* credits are available, regardless of the elapsed time since the last transmission.

The possible values for *both* attributes are 0x0 through 0xff. By default, i.e. if these attributes of a specific port *i* are not defined in the .INI file, then MAX_TIME_PORTi = 0x20 and MIN_TIME_PORTi = 0x8.

Example:

Suppose a user entered in the [Credits_Time] section the following two lines:

MIN_TIME_PORT*5* = 0x8

MAX_TIME_PORT*5* = 0x80

In this case, the shortest time period between two successive credit updates on a VL of port 5 will be 8 x 64 = 512 symbol-times; of course, this requires the availability of credits. The longest time period between two successive credit updates will be 128 x 512= 64K symbol-times. (0x80 = 12810)

7.2.26 [Speed5G]

This section is optionally used to configure Double Data Rate (DDR) speed for IB ports. The attributes used in this section are IBSpeed5_*i* and IBAutoNegSpeed5_*i*, where *i* is a port number (1,2,...24).

IBSpeed5_*i* is used to *force* the port speed to DDR. Possible values are Disable and Enable. See Section 7.2.26.1 for additional information.

IBAutoNegSpeed5_*i* is used to enable port speed auto-negotiation (SDR or DDR) to automatically set the fastest speed supported by both sides of the link. Possible values are Disable and Enable. Pay attention that a port forced to DDR speed by the IBSpeed parameter will not negotiate on the link speed even if its IBAutoNegSpeed5 parameter is enabled, but will try to use the DDR speed only.

Examples:

To force port 1 to DDR speed write:

IBSpeed5_1 = Enable

To configure port 2 to speed auto-negotiation mode write:

IBAutoNegSpeed5_2 = Enable

Under the following parameter setting, port 4 will attempt operating at DDR *without* auto-negotiation.

IBAutoNegSpeed5_4 = Disable
IBSpeed5_4 = Enable

7.2.26.1 Link Speed Lockup

An IB port *forced* to DDR (SerDes @ 5Gb/s) operation by means of the IBSpeed5_*i* parameter above may still be switched by the Subnet Manager to operate at SDR (SerDes @ 2.5Gb/s). To prevent this scenario, the user can enable the configuration parameter *LinkSpeedLockup*. This indicates to the SM that a device port forced to DDR operation will not switch to SDR operation even if asked to by the SM.

Possible values for the *LinkSpeedLockup* parameter are Disable (the default) and Enable.

Note: This mode of operation, where LinkSpeedLockup is enabled, serves as an extension to the functionlities defined in the *InfiniBand Architecture Specification (Release 1.2)*.

Example: Under the following parameters setting, port 4 will be able to operate at DDR only, and will not switch to SDR upon a request by the SM.

LinkSpeedLockup = Enable
IBSpeed5_4 = Enable

7.2.27 [VL_CAP]

This section is optionally used to set the

- maximum operational VLs for each port
- mode of credit allocation between port VLs

Maximum Operational VLs:

Each port can be configured to operate one of the following maximum number of VLs: max_1_vls (1 VL), max_2_vls (2 VLs), max_4_vls (4 VLs), or max_8_vls (8 VLs). By default, all ports are set to max_8_vls as follows:

PORT1 = max_8_vls

PORT2 = max_8_vls

...

PORT23 = max_8_vls

PORT24 = max_8_vls

Note that the actual number of Operational VLs per port is set by the Subnet Manager. For more details, see the section titled “Credit Allocation for Virtual Lanes” in the *MT47396 InfiniScale III Programmer’s Reference Manual, Revision 1.13* or later*, Document no. 2235PM*.

Credit Allocation Mode:

The mode of credit allocation is set using the parameter Cred_Alloc_Mode. If this parameter is 0 (default), then credit allocation complies with the *InfiniBand Architecture Specification*: credits for a single MTU (2KB) packet are allocated for each VL numbered above $2^{OperationalVLs-1}$ up to $2^{VLCap-1}$, and the remaining credits are allocated between the Operational VLs according to the InfiniScale III scheme of credit allocation. For more details, see the section titled “Credit Allocation for Virtual Lanes” in the *MT47396 InfiniScale III Programmer’s Reference Manual, Revision 1.13* or later*, Document no. 2235PM*.

If Cred_Alloc_Mode is 1, credits are allocated to the Operational VLs only according to the InfiniScale III scheme of credit allocation. VLs numbered above $2^{OperationalVLs-1}$ receive no credits at all.

Note: This mode (Cred_Alloc_Mode=1) is *not* compliant with the *InfiniBand Architecture Specification*.