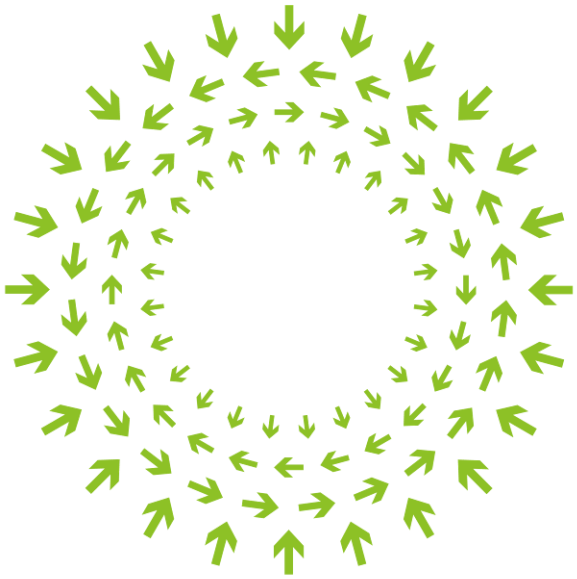

OPEN

Compute Project

Open Edge Switch

0.6

Author: Juha Miettinen, Nokia Solutions and Networks OY

Author: Samuli Toivola, Nokia Solutions and Networks OY

TABLE OF CONTENTS

License (OCP CLA Option)

Contributions to this Specification are made under the terms and conditions set forth in Open Compute Project Contribution License Agreement (“OCP CLA”) (“Contribution License”) by:

Nokia Solutions and Networks OY

Usage of this Specification is governed by the terms and conditions set forth in **Open Compute Project Hardware License – Permissive (“OCPHL Permissive”).**

Note: The following clarifications, which distinguish technology licensed in the Contribution License and/or Specification License from those technologies merely referenced (but not licensed), were accepted by the Incubation Committee of the OCP:

None

NOTWITHSTANDING THE FOREGOING LICENSES, THIS SPECIFICATION IS PROVIDED BY OCP "AS IS" AND OCP EXPRESSLY DISCLAIMS ANY WARRANTIES (EXPRESS, IMPLIED, OR OTHERWISE), INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, FITNESS FOR A PARTICULAR PURPOSE, OR TITLE, RELATED TO THE SPECIFICATION. NOTICE IS HEREBY GIVEN, THAT OTHER RIGHTS NOT GRANTED AS SET FORTH ABOVE, INCLUDING WITHOUT LIMITATION, RIGHTS OF THIRD PARTIES WHO DID NOT EXECUTE THE ABOVE LICENSES, MAY BE IMPLICATED BY THE IMPLEMENTATION OF OR COMPLIANCE WITH THIS SPECIFICATION. OCP IS NOT RESPONSIBLE FOR IDENTIFYING RIGHTS FOR WHICH A LICENSE MAY BE REQUIRED IN ORDER TO IMPLEMENT THIS SPECIFICATION. THE ENTIRE RISK AS TO IMPLEMENTING OR OTHERWISE USING THE SPECIFICATION IS ASSUMED BY YOU. IN NO EVENT WILL OCP BE LIABLE TO YOU FOR ANY MONETARY DAMAGES WITH RESPECT TO ANY CLAIMS RELATED TO, OR ARISING OUT OF YOUR USE OF THIS SPECIFICATION, INCLUDING BUT NOT LIMITED TO ANY LIABILITY FOR LOST PROFITS OR ANY CONSEQUENTIAL, INCIDENTAL, INDIRECT, SPECIAL OR PUNITIVE DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND EVEN IF OCP HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

1. Scope & Overview

Scope:

This document defines the technical specifications for the Open edge switch used in Open Compute Project.

Overview:

- Open edge switch is highly integrated, robust solution for Edge deployments:
- Supports L2/L3 Ethernet switch capabilities
- Supports indoors and outdoors deployment
- 18x SFP28 ports, Ethernet 10/25G
- 4x QSFP28 ports, supporting Ethernet 25/100G
- Compatibility with low-latency network requirements
- Open Edge compliant hardware, half width Open edge sled
- Compact footprint with 5* Switches in 3 RU or 3* Switches in 2 RU Open Edge chassis

Figure 1 Open edge switch in Open edge chassis

2. Product overview

Open Edge switch is Broadcom Monterey switch silicon-based switch with advanced synchronization capabilities.

Switch control processor is high performance Intel ATOM C3758, 8-core processor for extended feature set and application flexibility

Table 1 Open edge switch Feature List

Board Name	Open edge switch Mother Board
Form Factor	1U half width OE sled
Mother Board size	W 205.8mm x L 412.75mm, 14 layers, 1.8mm, 2 DIMMs
CPU	Intel Denverton C3758
Max Processor Wattage	25W
Switch Chip	Broadcom BCM5667x Monterey
Accelerator Board	Accelerator
Double oven oscillator module	Module Dimension: 42mmx130mm
Memory	DIMM Socket Count: 2 Total, CH A/B DIMM Types: DDR4 RDIMM, 1.2V, 2133/2400 DIMM Capacities: 8GB, 16GB
Front Side IO	(1) 18xEthernet/CPRI SFP28 25G (2) 4xEthernet QSFP28 100G (3) Synchronization in/out (2xHDMI) (4) Mini USB port for debugging (5) 3x System status LED (6) Synchronization status LED (7) Reset button
Network	LOM: Intel® Ethernet Controller I210 (I210)
Video	ASPEED AST2500 8MB DDR4 video memory
FAN	Dual rotor 4056 fan x4
ACPI	ACPI compliance, S0, S5 support. (* No S1 and S3 support.)
Power-Supply	Redundant power supply. 2x2000 Watt DC to DC power supply or 2x2000 Watt AC to DC power supply in chassis level.
TPM	TPM 2.0
Chassis	Recommended operating temperature: -5°C to 45°C Maximum operating temperature: -5°C to 55°C

	Non-operating temperature -40°C to 70°C Operating relative humidity 5% to 95% RH Non-operating relative humidity 5% to 95%RH

2.1 Block Diagram

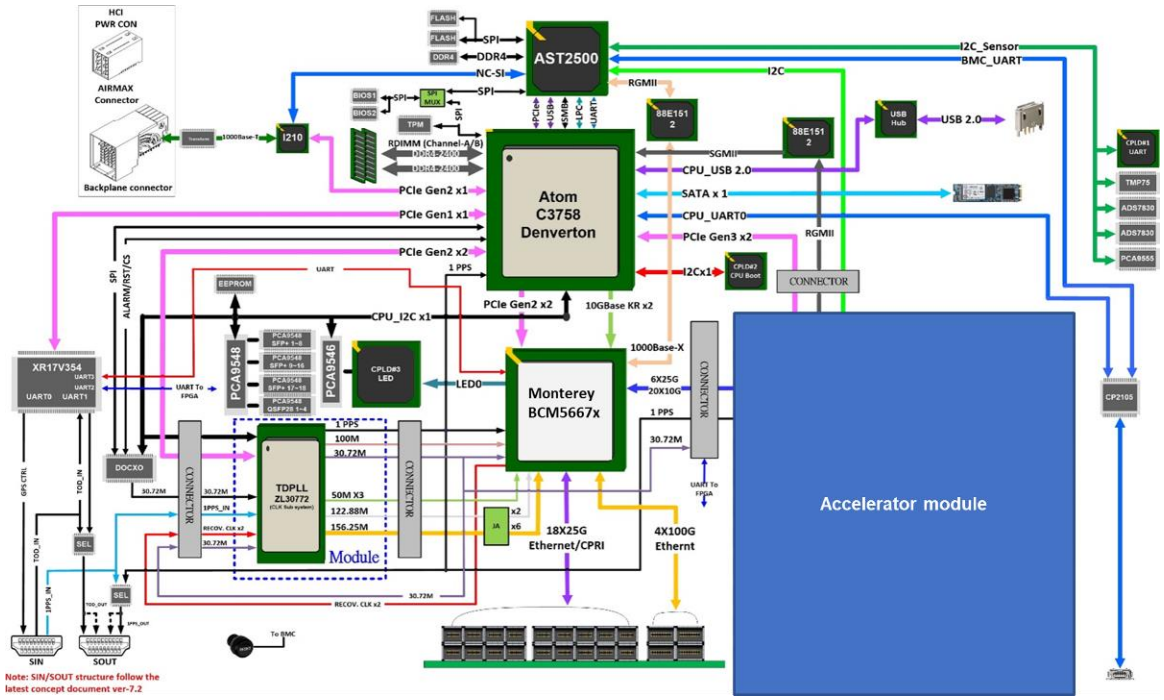

Figure 2 Open edge switch Board Block Diagram

2.2 Main Component Placement

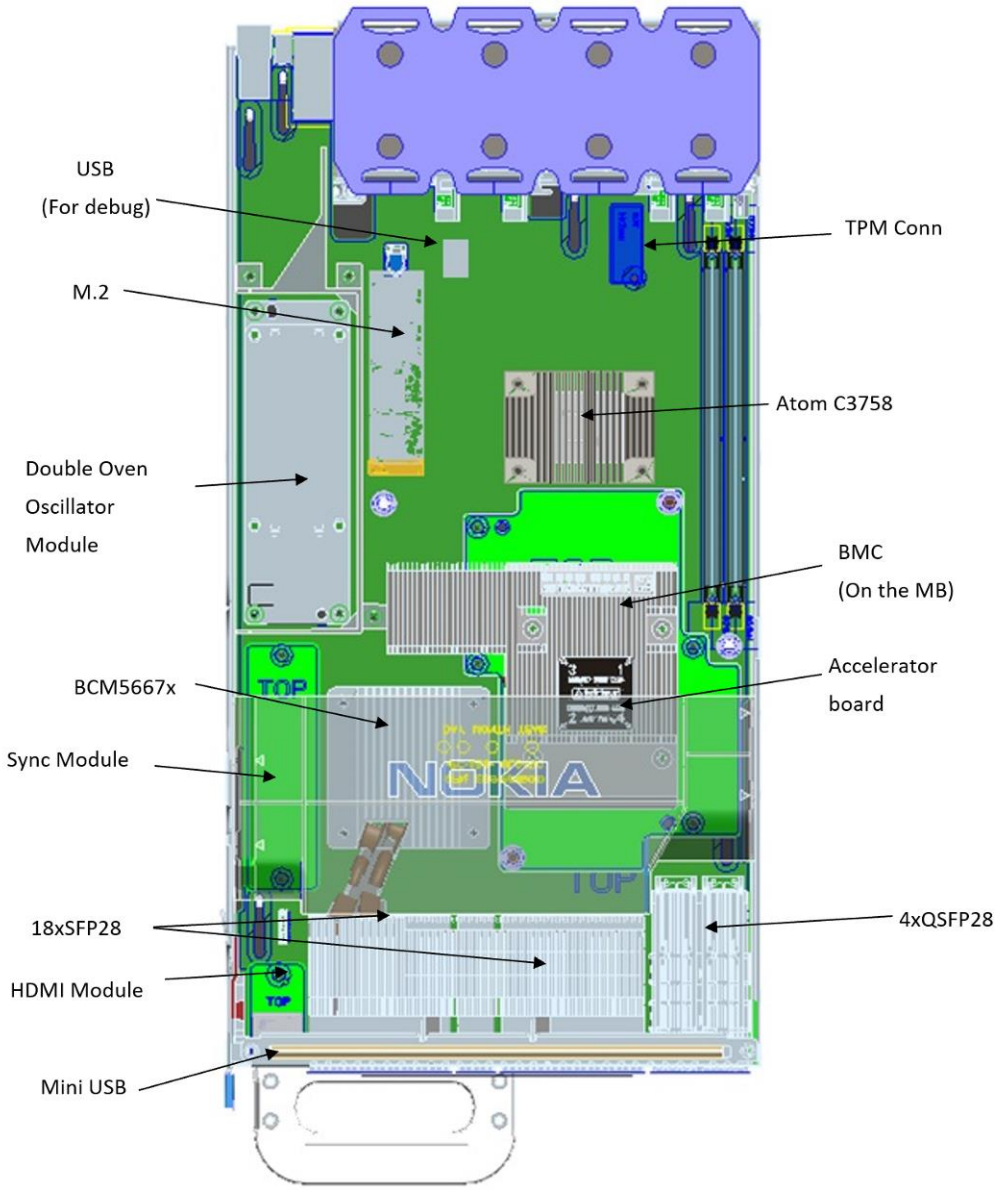

Figure 3 Open edge switch Board Placement

2.3 Form Factor

- Open edge switch is fully Open Edge chassis compatible hardware
- Open edge PSUs and RMC can be used together with switch
- Compact footprint with 5* Switches in 3 RU or 3* Switches in 2 RU Open Edge chassis

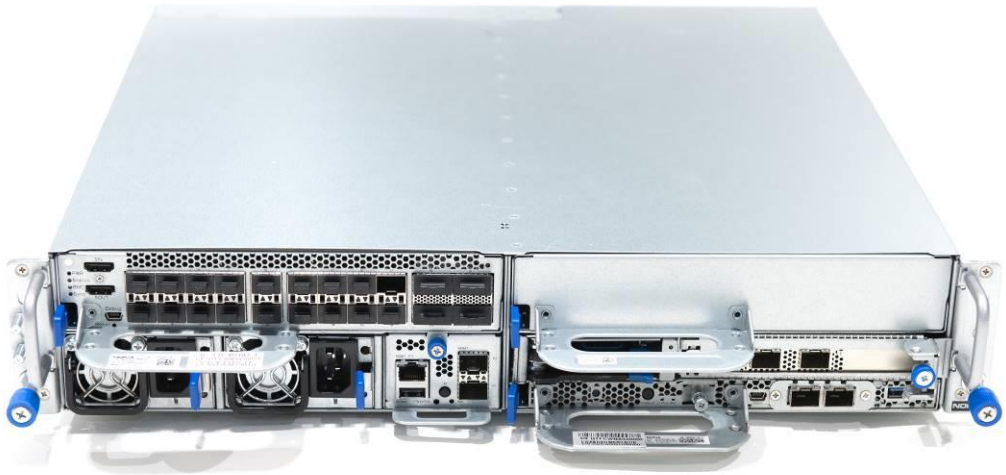

Figure 4 Open edge switch chassis placement

Figure 5 Open edge switch sled

Open edge switch main board dimension is 205.8mm x 412.75mm (WxL) as specified in Open Edge chassis base specification.

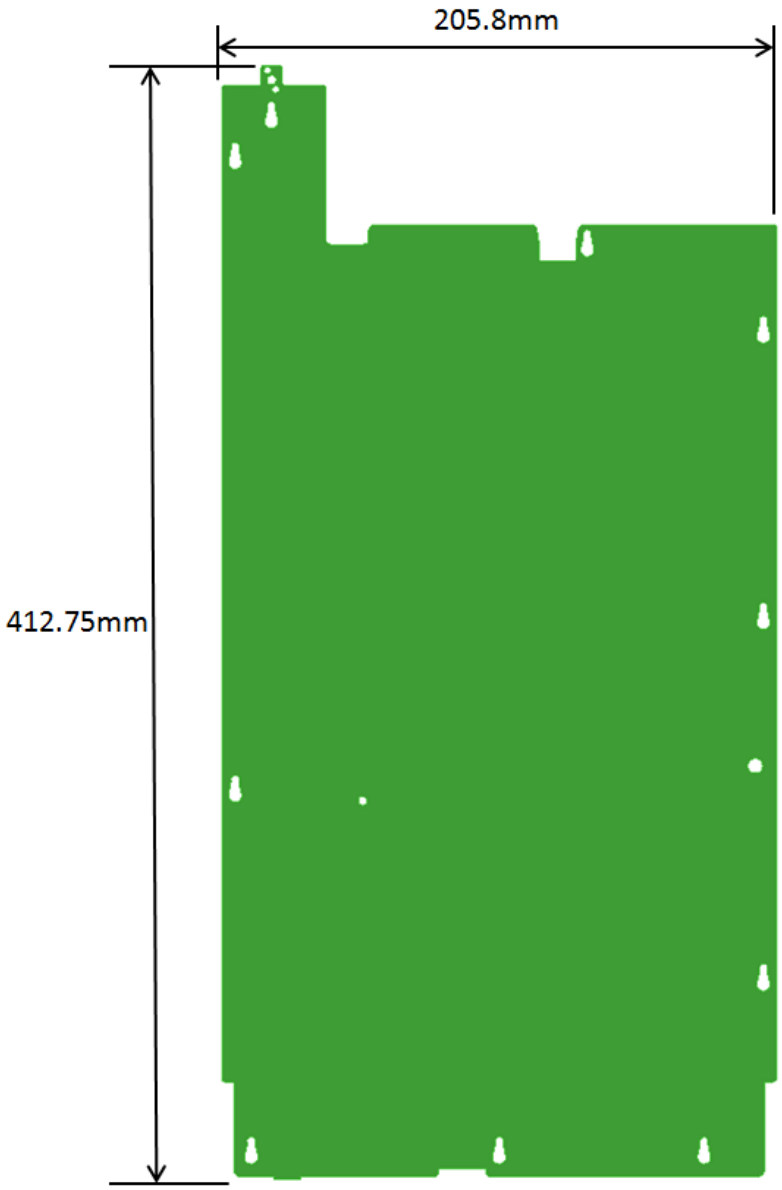

Figure 6 Open edge switch Mother Board Dimensions

2.4 CPU

Denverton adopts the Intel Atom C3758 processor. The Intel Atom Processor C3000 Product Family for Communications Infrastructure is the next generation of System-On-Chip (SoC) 64-bit processor built on the 14nm process technology and offers processor core frequencies from 1.5 GHz to 2.2 GHz. The

product family is ideal for energy efficient entry to midrange communications infrastructure solutions. This highly integrated SoC contains eight processor cores.

2.5 Memory

The Memory Controller is a dual channel DDR4 Memory Controller. The SoC supports a 40 or a 72 bit wide DDR4 interface per channel, 32 or 64 bits of data and 4 or 8 bits of ECC, and it does support both ECC and non-ECC DIMMs. RDIMM type are used in Open edge switch design.

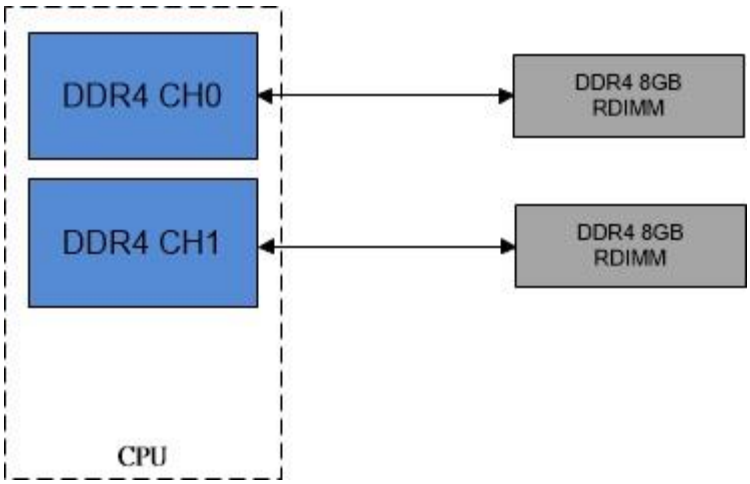

Figure 7 DRAM Port Configuration Options

Denverton uses 2*8GB DDR4 2400 DIMM SDRAM with ECC.

Table 2 Supported SDRAM Configurations

Memory Type	Voltage	DDR Standard	Memory Capacity	ECC
RDIMM	1.2V	DDR4-2400	2*8GB	Yes

2.6 PCIe

The SoC supports up to 16 PCI Express lanes that comply with the PCI Express Base Specification Revision 3.0. The maximum data rate is 8.0 GT/s. The 16 lanes can be configured to a maximum of eight, independent PCIe Root Ports (RP). The eight RPs are grouped in to two clusters, four RPs per cluster:

• Cluster 0 - RP0, RP1, RP2, RP3 are assigned Device Numbers 9, 10, 11, 12 decimals

• Cluster 1 - RP4, RP5, RP6, RP7 are assigned Device Numbers 14, 15, 16, 17

Twelve of the 16 PCIe lanes share their SoC signal pins with the integrated SATA0 and SATA1 controllers. Each of the 12 shared lanes must be configured as either PCIe or SATA through Soft Straps and BIOS.

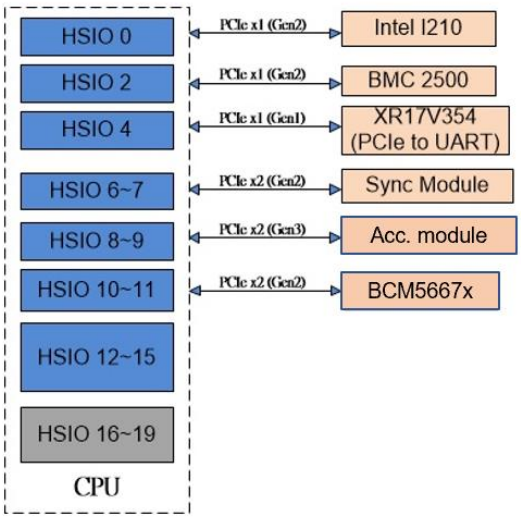

Figure 8 PCIe Topology Block Diagram

Table 3 PCIe ports connect from Denverton

	Connective
HSIO 0	Intel I210
HSIO 2	BMC AST2500
HSIO 4	XR17V354
HSIO 6~7	Sync module
HSIO 8~9	Accelerator board
HSIO 10~11	BCM5667x

2.7 PCB Stack-Up

STACKUP							
					Target Z (ohms) - MicroStrip		
					Target Z (ohms) - StripLine		
					Z tolerance		
					Z Type		
Layer#	**Material**	**Description**	**Copper Weight (oz)**	**Thickness (mil)**	**Tolerance (mil)**	**Glass Fabric**	**Er**
		Soldermask		0.60			3.8
1		TOP	0.5+plating	1.95			
	TU-883 Sp	PP		2.70	±0.709	1080x1	3.2
2		GND	0.5 (HVLP)	0.65			
	TU-883 Sp	CORE		4.00	±0.709	1035x2	3.2
3		IN1	0.5 (HVLP)	0.65			
	TU-883 Sp	PP		4.50	±0.709	2116x1	3.2
4		GND1	0.5 (HVLP)	0.65			
	TU-883 Sp	CORE		4.00	±0.709	1035x2	3.2
5		IN2	0.5 (HVLP)	0.65			
	TU-883 Sp	PP		4.50	±0.709	2116x1	3.2
6		GND2	1 (HVLP)	1.30			
	TU-883 Sp	CORE		4.00	±0.709	1035x2	3.2
7		VCC	2 (RTF)	2.60			
	TU-883 Sp	PP		4.00	±0.709	106x2	3.2
8		VCC1	2 (RTF)	2.60			
	TU-883 Sp	CORE		4.00	±0.709	1035x2	3.2
9		GND3	1 (HVLP)	1.30			
	TU-883 Sp	PP		4.50	±0.709	2116x1	3.2
10		IN3	0.5 (HVLP)	0.65			
	TU-883 Sp	CORE		4.00	±0.709	1035x2	3.2
11		GND4	0.5 (HVLP)	0.65			
	TU-883 Sp	PP		4.50	±0.709	2116x1	3.2
12		IN4	0.5 (HVLP)	0.65			
	TU-883 Sp	CORE		4.00	±0.709	1035x2	3.2
13		GND5	0.5 (HVLP)	0.65			
	TU-883 Sp	PP		2.70	±0.709	1080x1	3.2
14		BOTTOM	0.5+plating	1.95			
		Soldermask		0.60			3.8
			Total	**69.50**	**±10%**		

Figure 9 Open edge switch Mother Board Stackup

3. BIOS

The BIOS has capability to upgrade the BIOS flash from Host CPU or BMC. When upgrade BIOS flash from BMC, it can be done in background without interrupting any functionality of host.

4. BMC

The Open edge switch's Board Management Controller (**AST2500**) is a highly integrated single-chip solution, integrating several devices typically found on servers.

Video-over-IP

• Video Redirection up to 1920x1200 resolution. YUV444/YUV420 Video Compression, 24 bits video compression quality

DDR3L/DDR4 SDRAM Controller

• Support external 16-bit DDR3L/DDR4 SDRAM data bus width

• Maximum memory clock frequency

– DDR3L: 800MHz (DDR3-1600)

– DDR4: 800MHz (DDR4-1600)

GPIO Controller

• Directly connected to APB bus

• Support up to 228 GPIO pins, which are 29 sets

• Each GPIO sets can be programmed to accept command from ARM, LPC(SIO), or Coprocessor.

• Programmable output mode: Push-Pull or Open-Drain

• Some GPIOs support Schmitt type input buffer for noise immunity

• 4 out of the 228 GPIO pins are with 16mA driving strength, others are 8mA driving strength

• 16 out of the 228 GPIO pins that can support 1.8V mode.

• Support 8 sets of GPIO pass through (1 GPIO IN -- > 1 GPIO OUT) pin with internal switch control, it is useful for some button function control.

4.1 Management Network Interface

The LOM solution in Open edge switch is Intel Ethernet controller I210-AT. I210-AT is a single port, compact, low power component that supports GbE designs. The I210 offers a fully integrated GbE Media Access Control (MAC). The I210-AT enables 1000BASE-T implementations using an integrated PHY. Below are the features of I210-AT

- Operating Temperature: 0 to 70 °C.
- PCIE v2.1 (2.5 GT/s) x1 is used by the I210 as a host interface. In Open edge switch, the PCIE interface connects to CPU PCIE port.
- Network Interface: 1000 Base-T. In Open edge switch, the 1000 Base-T is routing on the server sled towards back plane →PDB→ finally to the Broadcom Ethernet switch BCM53158 on RMC as shown in Figure 8.

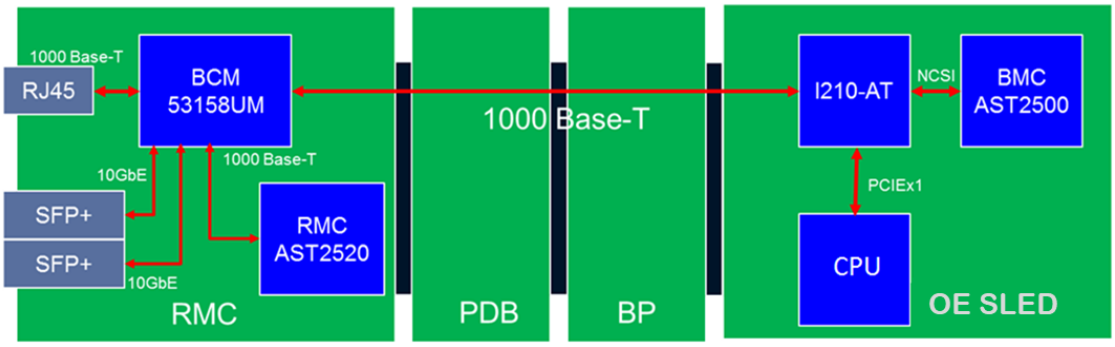

Figure 10 1000 Base-T routing topology

It provides the below capabilities for the BMC on the Open Edge switch sled:

1. A dedicated 1000 BASE-T Ethernet port for hardware management. User can remote manage the server sled via this interface.
2. RMC can get the sled information via this interface.

➢ Manageability: NCSI interface. This is accomplished by providing mechanisms by which manageability network traffic can be routed to and from BMC.
Pass-Through (PT) is the term used when referring to the process of sending and receiving Ethernet traffic over the sideband interface. The I210 has the ability to route Ethernet traffic to the host operating system as well as the ability to send Ethernet traffic over the sideband interface to an external BMC. The I210 supports two sideband interfaces:

- SMBus
- NC-SI

Only one mode of sideband can be active at any given time.

In Open edge switch, we used the NCSI interface to connect BMC. The usable bandwidth for either direction is up to and 100 Mb/s for the NCSI interface.

➢ Advanced Features:

- Jumbo frames
- Interrupt moderation, VLAN support, IP checksum offload.
- PCIE OBFF (Optimized Buffer Flush/Fill) for improved system power management.
- Four transmit and four receive queues.
- RSS and MSI-X to lower CPU utilization in multi-core systems.
- ECC – error correcting memory in packet buffers.

4.2 Local Serial Console

Denverton CPU have two UART port, it provided serial communication capabilities, which allow communication with the modem or other external devices using RS-232 protocol. Each UART port is separately configured and can be run as an asynchronous link from 1200 baud to 4 M baud.

The UART0 is used debugging/management, it is connected to front panel to mini USB connector through the CP2105(UART to USB converter). The UART1 is connect to BMC.

Open edge switch provides Host and BMC UART for development purpose. Both UARTs connect to a UART to USB converter on the board. The USB connect to a mini USB connector at the front IO. User can an USB cable to link the USB port to laptop to access the UARTs.

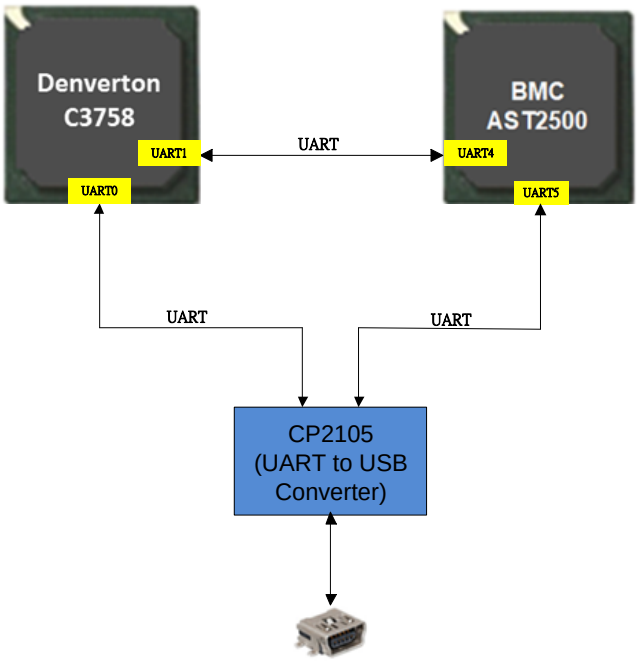

Figure 11 UART interface on Denverton

5. Monterey Switch

The Broadcom BCM5667x family is a set of new high-performance, Radio-over-Ethernet (RoE) Switches optimized for advanced mobile networking applications. BCM5667x allow direct connection to cellular radios using Common Public Radio Interface (CPRI) links and act as a bridge between the traditional mobile network and Ethernet infrastructure of modern, cloud-based platforms.

5.1 TSCe (PM 4x10) SerDes

There are up to 6 TSCe instances in the BCM5667x device. Each TSCe contains four SerDes lanes. Each lane can be configured as a single port or multiple lances aggregated into a single port. A single TSCe consists of four SerDes Lanes. Each lane can operate between 1.25 Gb/s, 3.125 Gb/s, 6.25 Gb/s, 6.5625Gb/s, 10.3125Gb/s and 10.9375 Gb/s. 10G BASE-KR is supported in a single lane. In this configuration, each TSCe instance can support four KR ports.

5.2 TSCe (PM 4x10) SerDes

The TSCf SerDes is the versatile physical layer interface for the BCM5667x, specifically designed to support up to 100Gb/s. It supports CPRI line rates ranging from 2.5 Gb/s through 24.3 Gb/s per serial link. It also supports Ethernet line rates ranging from 10.3125 Gb/s through 27.34375 Gb/s per serial link.

There are a total of 10 TSCf instances, which consist of 4 x TSCf(PM) cores and 6 x TSCf(CPM) cores in the BCM5667x. Each TSCf contains four SerDes lanes. Each lane can be configured as a signal port or multiple lances and be aggregated into a signal port.

Open edge switch system uses 18 Ethernet/CPRI interfaces and 4 x 100G interfaces in front panel. There are also 20 x 10Gb/s and 6 x 25Gb/s connect to accelerator board which will be defined in the future. There are 2 10GBase-KR connect to CPU and one 1000Base-X connect to BMC.

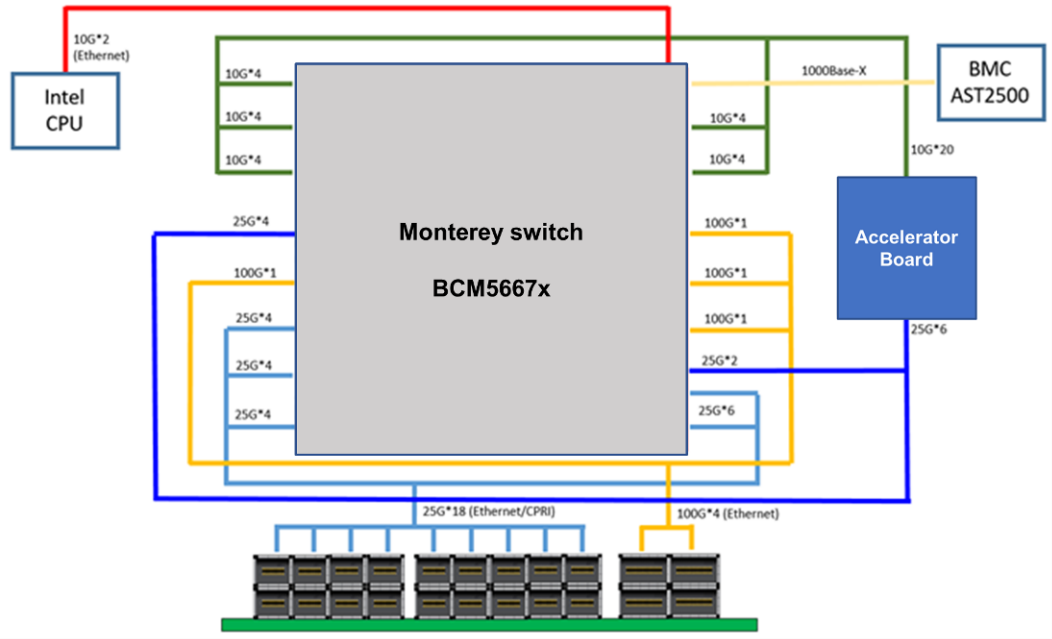

Figure 12 BCM5667x High-speed I/O mapping

6. Rear Side Power, I/O and Midplane
    6.1 Overview of Footprint and Population Options
    6.2 Rear Side Connectors

7. Mechanical
    7.1 Sled Dimensions

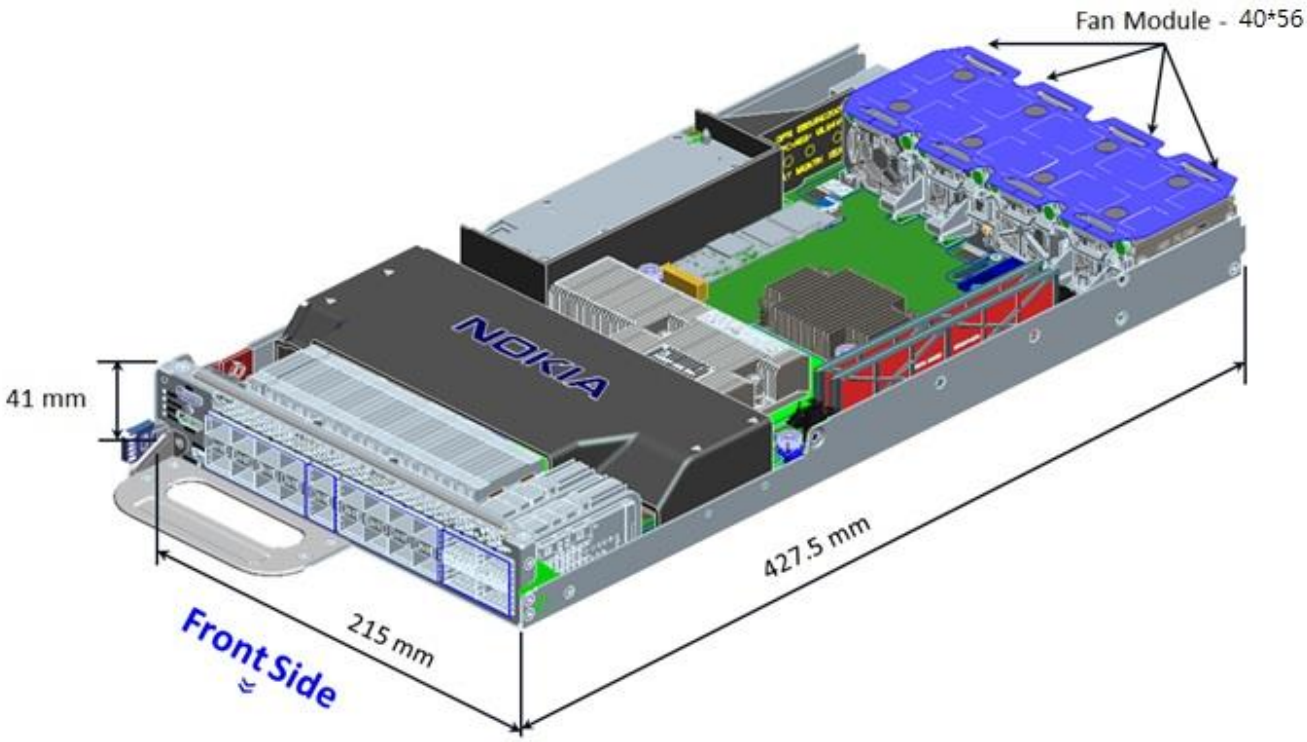

Figure 13 Open edge switch System Dimensions

7.1 Front panel interfaces

Switch provides optimal capacity in small footprint:

- 18x SFP28 ports, Ethernet 10/25G
- 4x QSFP28 ports, supporting Ethernet 25/100G
- In addition, switch front panel provides connectors for synchronization input and output (1PPS/ToD)

•

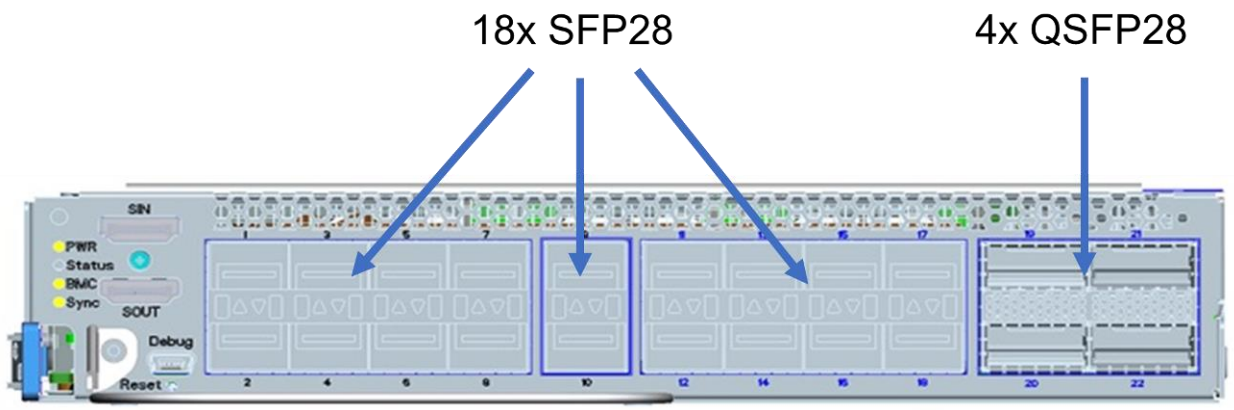

Figure 14 Open edge switch Front Panel Connectors

7.2 Leds and indicators

Figure 15 Open edge switch Front Panel LED

Table 4 Power / ID led

Name	Color	Condition	Behavior
PWR/ID LED	Blue	Off	System id off
		Blinking	Identify the system
		Solid On	System is on

Table 5 System Status Led

Name	Color	Condition	Behavior
STATUS LED	Green	On	Steady Green Color when system is operating normally
	Amber	Blinking	Blinking amber color when system have abnormal event

Name	Color	Condition	Behavior
BMC Heartbeat LED	Green	On/Off	BMC is not ready
		Blinking	BMC is ready

Name	Color	Condition	Behavior
Synchronization Status LED	Green	On	Synchronization is ready
		Blinking	Synchronization is on going
	Red	On	Synchronization is Loss

7.3 Ethernet port Leds

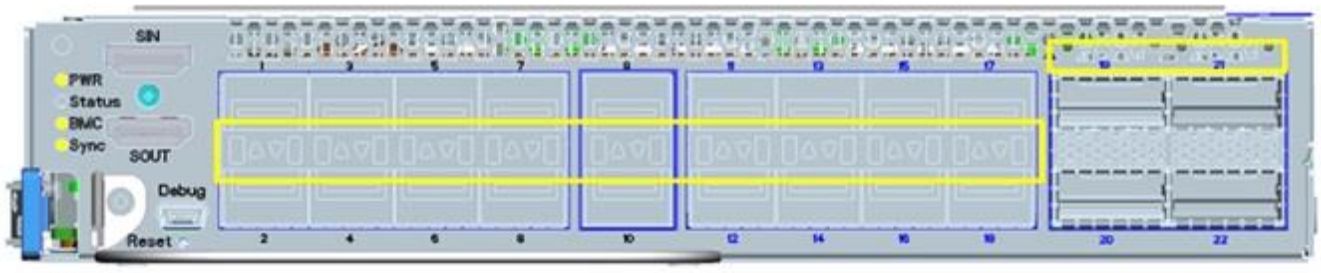

Figure 16 Open edge switch SFP29 and QSFP28 Led locations

SFP28 Port LED

Each SFP28 port has LED to indicate port status. The LED definition as below:

Table 6 SFP28 Port Status LED

Mode	Link	Active	No Link	Fault
25G	Green	Green blink	Off	Off
10G	Amber	Amber blink	Off	Off
1G	Amber	Amber blink	Off	Off

QSFP28 Port LED

Each QSFP28 port has LED to indicate port status. The LED definition as below:

Table 7 QSFP28 Port Status LED

Mode	Link	Active	No Link	Fault
100G	Green	Green blink	Off	Off
50G	Amber	Amber blink	Off	Off
40G	Amber	Amber blink	Off	Off
25G	Amber	Amber blink	Off	Off
10G	Amber	Amber blink	Off	Off
1G	Amber	Amber blink	Off	Off

7.4 Heat Sinks and ILM

MAC Heat Sink

- Fin Dimension=134.2*40*6.7 mm^3 (+/-2 mm each)

- Material: AL base + AL Fin with two heat pipes.

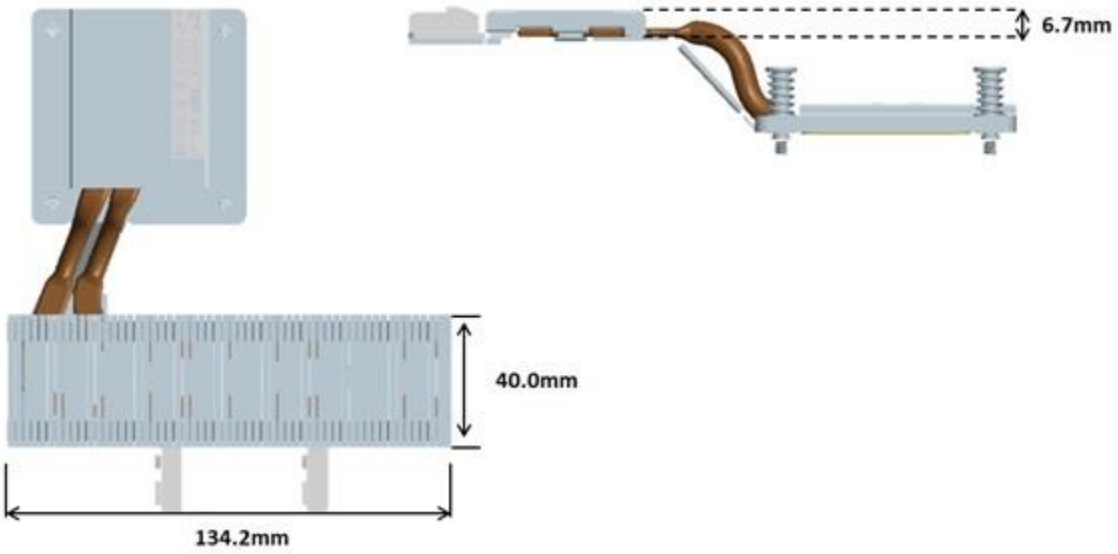

Figure 17 Monterey and SFP28 heat sink

CPU Heat Sink

- Dimension=60*45*18 mm^3 (+/-2 mm each)
- Material: AL6063

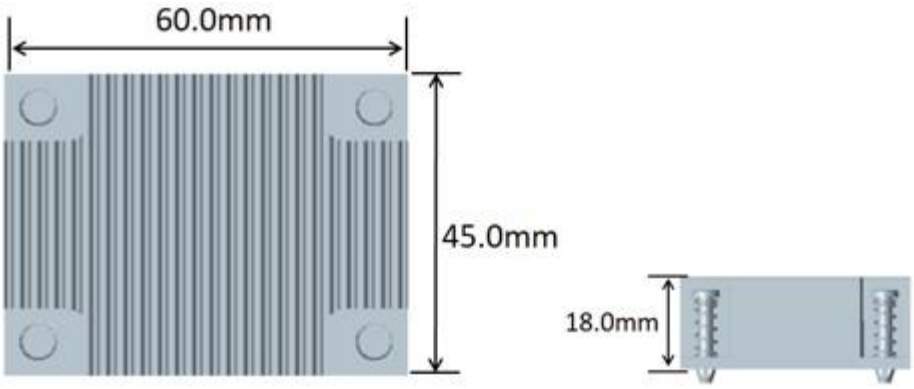

Figure 18 CPU heat sink

Accelerator Heat Sink

- Dimension=116.35*99*20.4 mm^3 (+/-2 mm each)

- Material: AL base + Cu Fin with two heat pipes.

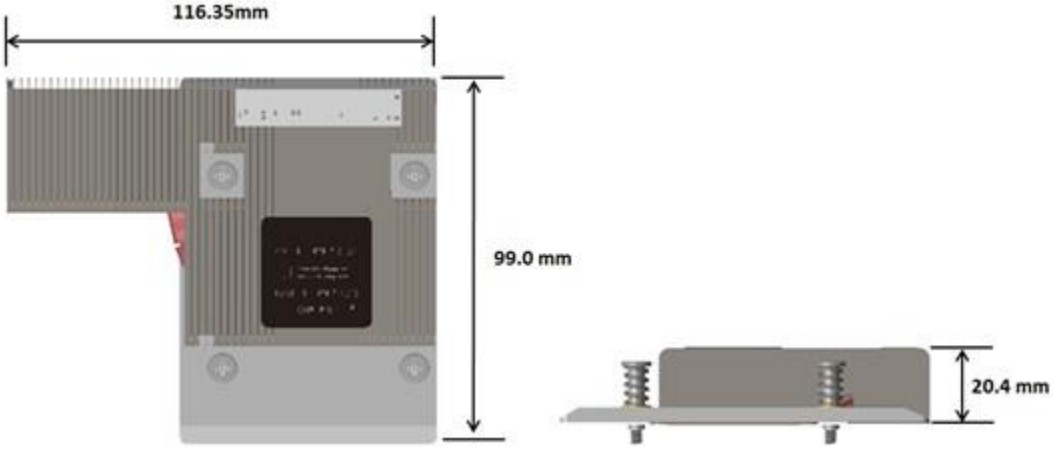

Figure 19 Accelerator heat sink

7.5 Fan Design

Fan x4

- Dimension=4056 (Dual rotor)
- Voltage=12 V
- PWM Frequency= 25 KHz

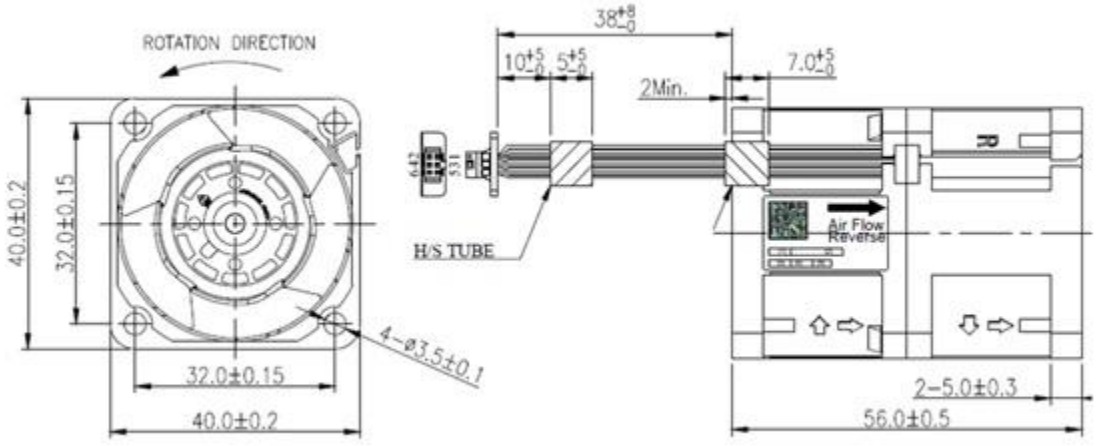

Figure 20 4056 Fan Drawing

7.6 Operating Environment

Thermal Environment/Support Condition

Unit shall comply with all applicable requirements of GR-63-CORE (Section 4.1) as follows:

- Recommended operating temperature: +5 C to +45 C
- Maximum operating temperature: -5 C to +55 C
- Non-operating temperature: -40 C to +70 C
- Operating humidity: 5% to 95%
- Non-operating humidity: 5% to 95%
- System startup temperature min +5 C.

8. Appendix

8.1 User Guidance

In the industry, there is a known vulnerability, CVE-2019-6260, when using ASPEED AST2500.

There are some approaches to mitigate:

a. Designers can refer to AST_usrGuide_QuickRef for CVE-2019-6260 to fine tune the firmware.

b. Designers can use another BMC chip to replace ASPEED AST2500 in the derivative designs.

8.2 Nokia actions to Aspeed QuickRef

According to CVE-2019-6260, ASPEED AST2400 and AST2500 Baseband Management Controller (BMC) hardware and firmware implement Advanced High-performance Bus (AHB) bridges, which allow arbitrary read and write access to the BMC's physical address space from the host. The LPC, PCIe and UART AHB bridges are all explicitly features of ASPEED's designs for recovering BMC FW during firmware development or to allow the host to drive the BMC hardware even without any firmware on BMC. The CVE applies to the Eight (8) specific cases of iLPC2AHB bridge Pt I, iLPC2AHB bridge Pt II, PCIe VGA P2A bridge, DMA from/to arbitrary BMC memory via X-DMA, UART-based SoC Debug interface, LPC2AHB bridge, PCIe BMC P2A bridge, and Watchdog setup.

[Suggested Mitigation of CVE]

UART-based SoC Debug interface;

Nokia AirFrame server debugging interface (mini-USB) shall be used by an authorized person only.

Users shall limit the access to management interfaces which poses vulnerability risk. As AHB is board internal bus the AHB security issue is inside the server, so the security risk is fairly small. The board’s external interfaces are securely protected in design.

Nokia recommends to use ASPEED AST2600 which provides the vulnerability correction and using SecureBoot to enhance security.