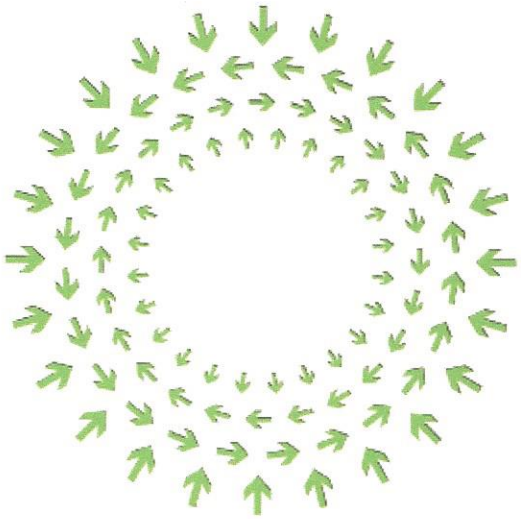

OPEN

Compute Project

Broadcom Open 1.0

Leaf and Spine

Switch Specification

October 2013

Revision History

Revision	Date	Author	Changes
01	March 18, 2013	Ben Askarinam	Created
02	April 7th, 2013	Ben Askarinam	General updates
03	April 8th, 2013	Ben Askarinam	Updated Figures 14 1nd 15
04	April 14th, 2013	Ben Askarinam	Added section 9
05	April 15th, 2013	Ben Askarinam	Misc. Update
06	April 18th, 2013	Ben Askarinam	Added board outlines
07	May 6, 2013	Fazil Osman	Incorporated Sujal's changes
08	May 9, 2013	Sujal Das	Added configuration options
09	May 10, 2013	Sujal Das	Miscellaneous Update
10	May 11th, 2013	Ben Askarinam	Updated switch section
11	May 13th, 2013	Ben Askarinam	Updated figures 13, 14, 15, 16 & 17
12	May 17th 2013	Sujal Das	Reduced the size of the document by re-embedding the figures 13-17
13	June 14th, 2013	Ben Askarinam	Updated Section 6 a) b) Table 4 c) Figure 8 d) Figure 9
14	June 27th, 2013	Sujal Das	Changed document format. Added more detailed descriptions to many sections. Removed vendor proprietary information.
15	July 15th, 2013	Ben Askarinam	Updated figures and sections 11-18
16	September 11th, 2013	Prem Jonnalagadda	Replaced SAAP with CAP in clear text. Fixed the Diagrams also to replace SAAP with CAP.
17	September 30th, 2013	Ben Askarinam	Added TIIS Option A & Option B Added Figure 3 & Figure 4 for 32 Ports of 40GbE.
18	October 22nd, 2013	Heather Meeker on behalf of OCP and Broadcom	Added Section 3, Licensing terms for Semiconductor Devices
19	October 29th, and 30th, 2013	Sujal Das, Ben Askarinam, Prem Jonnalagadda	Technical and Formatting Updates.

Table 1, Document Revision

Contents

TABLE OF FIGURES

TABLE OF Tables

This specification is entered into in conjunction with and incorporated by reference in its entirety into the Open Web Foundation Contributor License Agreement executed by Broadcom Corporation on [_____], 2013 (CLA).

1 Scope

This document defines the technical specifications for the Broadcom Open Switch Platform that will be used as a 10GbE Leaf and 40GbE Spine Switch in an Open Compute Project. The Leaf Switch may also be referred to as Top of Rack (ToR) Switch and the Spine Switch may be referred to as an Aggregation Switch. Such Leaf and Spine Switches can be used to create CLOS-based fat and flat topologies for building high performance network infrastructures.

2 Overview

This document describes the technical specifications for the Broadcom Open Switch Platforms that will be used as a 10GbE/40GbE Top of Rack (ToR) Switch and 40GbE Aggregation Switch in an Open Compute Project. These switches are designed to be compatible with the OCP Open Rack specification.

The Open Switch Platform fits in the Open Rack 1U form factor and is optimized for high capacity switching and enabling cloud, Web 2.0, enterprise and service provider applications. By enabling use of embedded as well as server class standard Linux distributions as the foundation for network operating systems that can run on this platform, a plethora of software defined networking (SDN) applications such as network automation, network monitoring, network instrumentation , network virtualization as well as specialized network functions can be deployed on this platform. Modular pluggable options for various CPUs make this platform highly flexible.

The Open Switch Platform is optimized for cost and power and meets the requirements for Leaf and Spine Switches that can be used by large data center, cloud computing and telecom service providers. It has been designed to meet the small form 1U form factor to maximize the space in the rack for

compute and storage. It has also been designed to provide the flexibility necessary to meet the perceived use cases for Leaf and Spine Switches for the foreseeable future. Moreover, the core switch module can be repackaged to fit in a standard Enterprise 19 inch rack increasing the market available to the manufacturers of this Open Switch Platform. The Enterprise version is not described in this document. This Platform will be used to deploy 10GbE-based Compute servers and Storage servers in the datacenter with the flexibility to connect into either a 10G or 40G aggregation layer switches.

Like previous Open Compute projects, this Open Switch Platform in its basic configuration is power-optimized and barebones, designed to provide the lowest capital and operating costs. With optional pluggable modules, the Open Switch Platform enables extensibility to support emerging SDN applications. The Switch Platform enables customer driven (versus OEM or ODM driven) hardware core features and requirements and will drive an open standards approach that should generate volume and competitive pricing.

3 License

All semiconductor devices that may be referred to in this specification, or required to manufacture products described in this specification, will be considered references only, and no intellectual property rights embodied in or covering such semiconductor devices shall be licensed as a result of this specification or such references. Notwithstanding anything to the contrary in the CLA, the licenses set forth therein do not apply to intellectually property rights included in or related to the semi-conductor devices identified in the Specification. These references include without limitation the references to devices listed below. For clarity, no patent claim that reads on such semiconductor devices will be considered a "Granted Claim" under the applicable Contributor License Agreement for this Specification.

Broadcom StrataXGS Trident II switch
AMD x86 CPU
Broadcom XLP432 communication processor
10/40GbE Switch BCM56850
Broadcom SERDES Repeater BCM84328
Altera CPLD EPM2210F324C5N
AMD X86 CPU T40N
AMD IO HubA55E FCH
Micron Technology SODIMM Memory MT18KSF25672HZ
Broadcom GE MAC/PHY Controller BCM5725
Broadcom 32 Core MIPS CPU XLP432
Micron Technology SODIMM Memory MT18KSF25672HZ
Altera CPLD EPM2210F324C5N

4 Product Descriptions and Requirements

Figure 1 & Figure 2 show the block diagram for 10GbE/40GbE model of the 19 inch rack version & Open Rack version.

Figure 3 & Figure 4 show the block diagram for the 32 ports of 40GbE model of the 19 inch rack version & Open Rack version

It supports the following high density switch configurations:

a) 48x 10GbE + 2x 40GbE
b) 48x 10GbE + 4x 40GbE
c) 48x 10GbE + 6x 40GbE
d) 48x 10GbE + 12x 40GbE
e) 60x 10GbE + 2x 40GbE
f) 60x 10GbE + 4x 40GbE
g) 60x 10GbE + 6x 40GbE
h) 72x 10GbE
i) 96x10GbE[1]
j) 32 x 40GbE[2]

[1] Requires TIIS Option B
[2] Requires TIIS Option B

The Switch Platform consists of the following modules:

1. A Trident II Switch module (TIIS) –This module consists of a Broadcom StrataXGS® Trident II switch chip with 1.28Tbps of bandwidth.
2. An Control and Application Processor Module (CAP) - The CAP module uses a low power AMD x86 CPU by default to provide the following functions:
   a) It runs the Broadcom SDK switch software to program and configure the Trident II switch.
   b) It provides an environment where datacenter operators can develop their own applications in a software development environment familiar to them such as Ubuntu Linux. Such applications can be used to automate the deployment and management of the switches as well as compute and storage servers in the rack.

   It is envisioned that many different CAP module options will be available on designs based on the Open Switch Platform specification to meet the application and cost requirements of end users. For example, Broadcom XLP series based MIPS or FreeScale PPC based CAP modules may be provided.
3. A Data Plane Processor Module (DPP) – This optional module is used to provide acceleration to a variety of data plane functions such as encryption and decryption, NAT and other sophisticated packet inspection and modification capabilities. The BRCM XLP432 communication processor is used

on the DPP module enabling specialized network functions on the Open Switch Platform.

4. IO Modules – The pluggable IO Modules provide the most flexibility in port configuration to fit the cost, rack size, power, wiring and network topology design needs of the end user. There are two kinds of IO Modules – SFP+ for 10GbE ports and QSFP for 40GbE ports.

   The Open Leaf Switch Platform has been designed to support up to 14 such IO modules. A maximum of 12 of the IO slots can be populated with the Quad SFP+ 10GbE IO modules to provide a maximum of 48 10GbE ports. A maximum of 2 of the IO slots can be populated with the Hex QSFP 40GbE IO modules to provide a maximum of 12 40GbE ports.

   The Open Spine Switch Platform has been designed to support up to 6 QSFP IO modules. A maximum of 5 of the IO slots can be populated with the Hex QSFP 40GbE IO modules and 1 of the IO slots can be populated with Dual QSFP+ 40GbE IO module to provide a maximum of 32 40GbE ports.

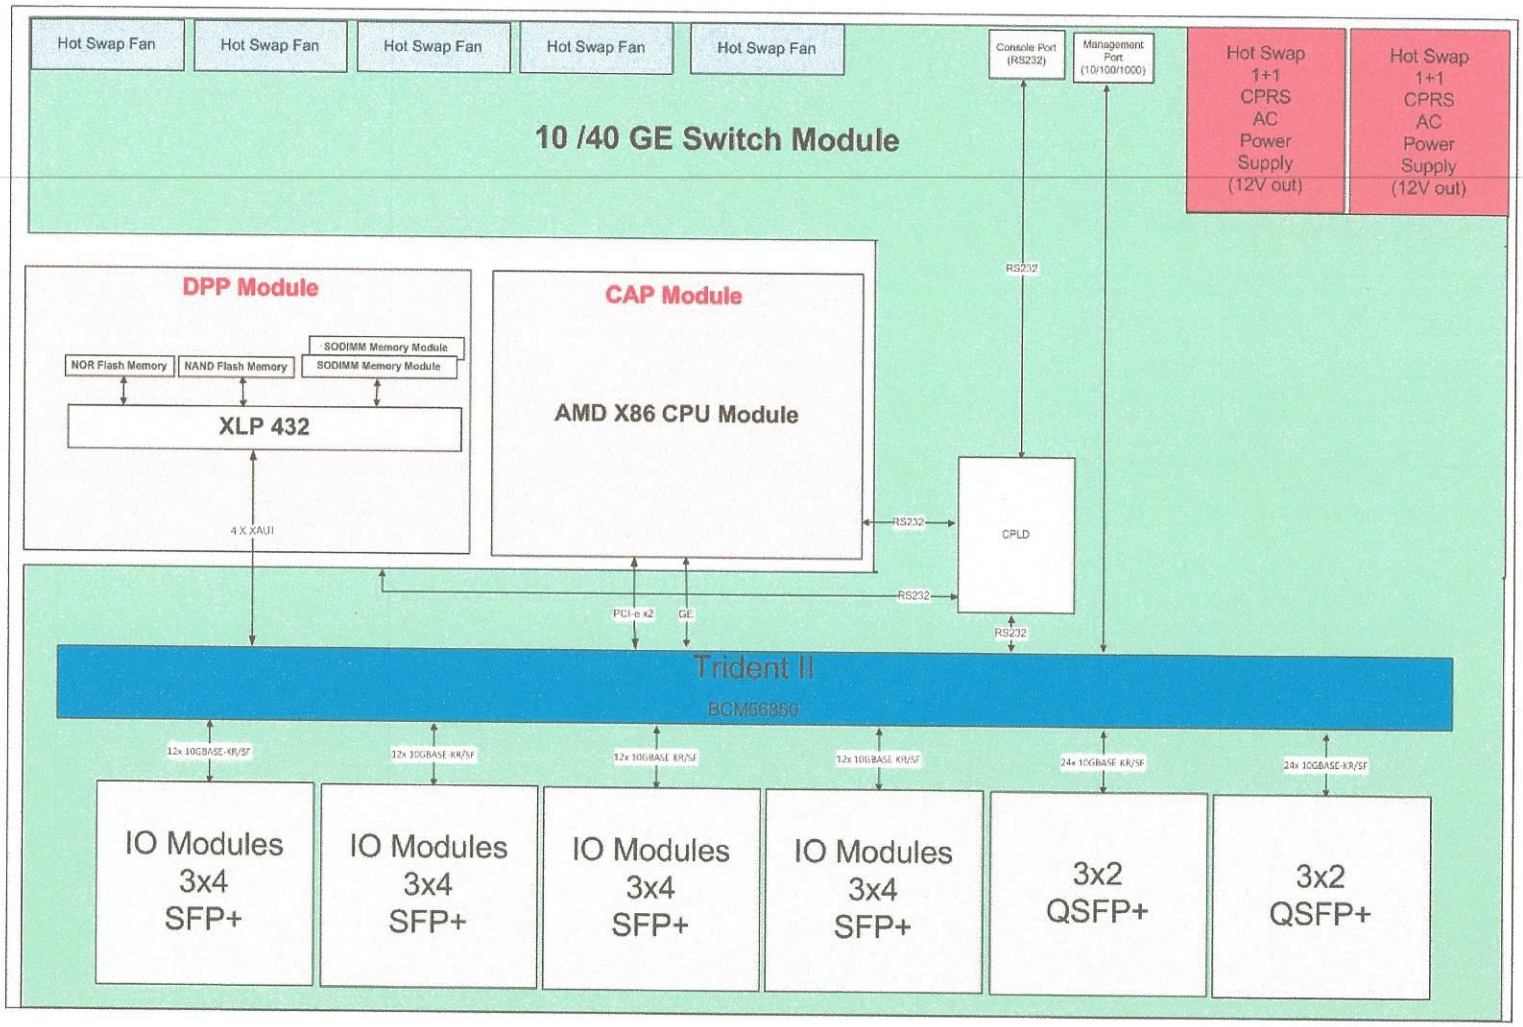

Figure 1 10GbE / 40GbE OCP Open Switch Platform Block Diagram for Standard 19 Inch Rack

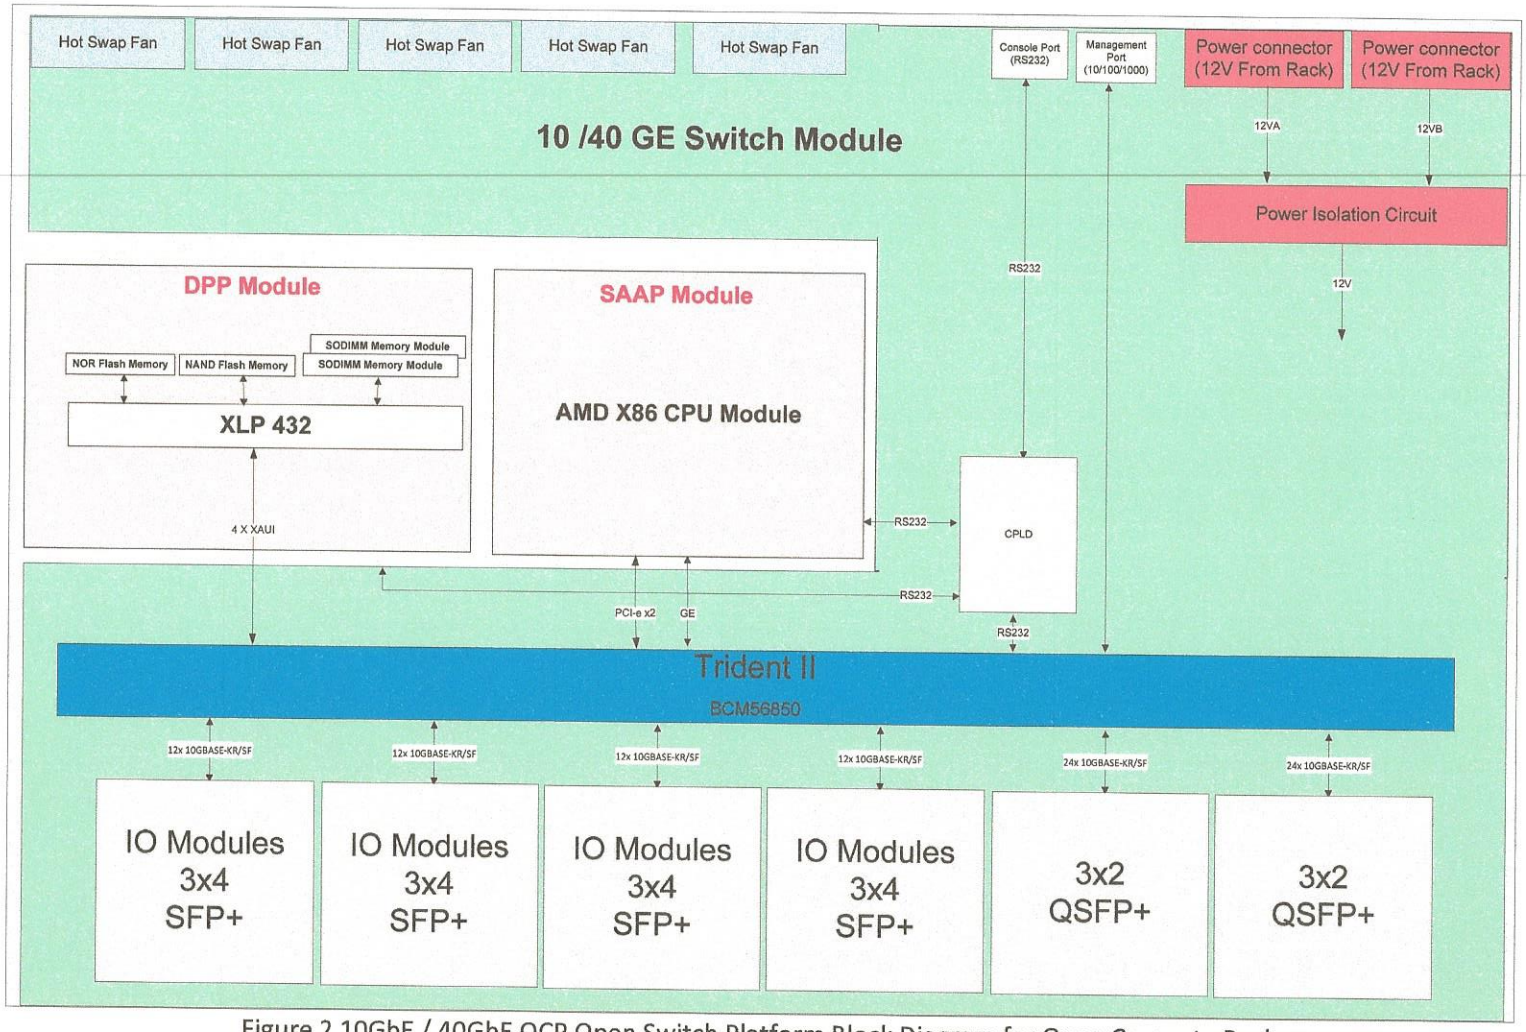

Figure 2 10GbE / 40GbE OCP Open Switch Platform Block Diagram for Open Compute Rack

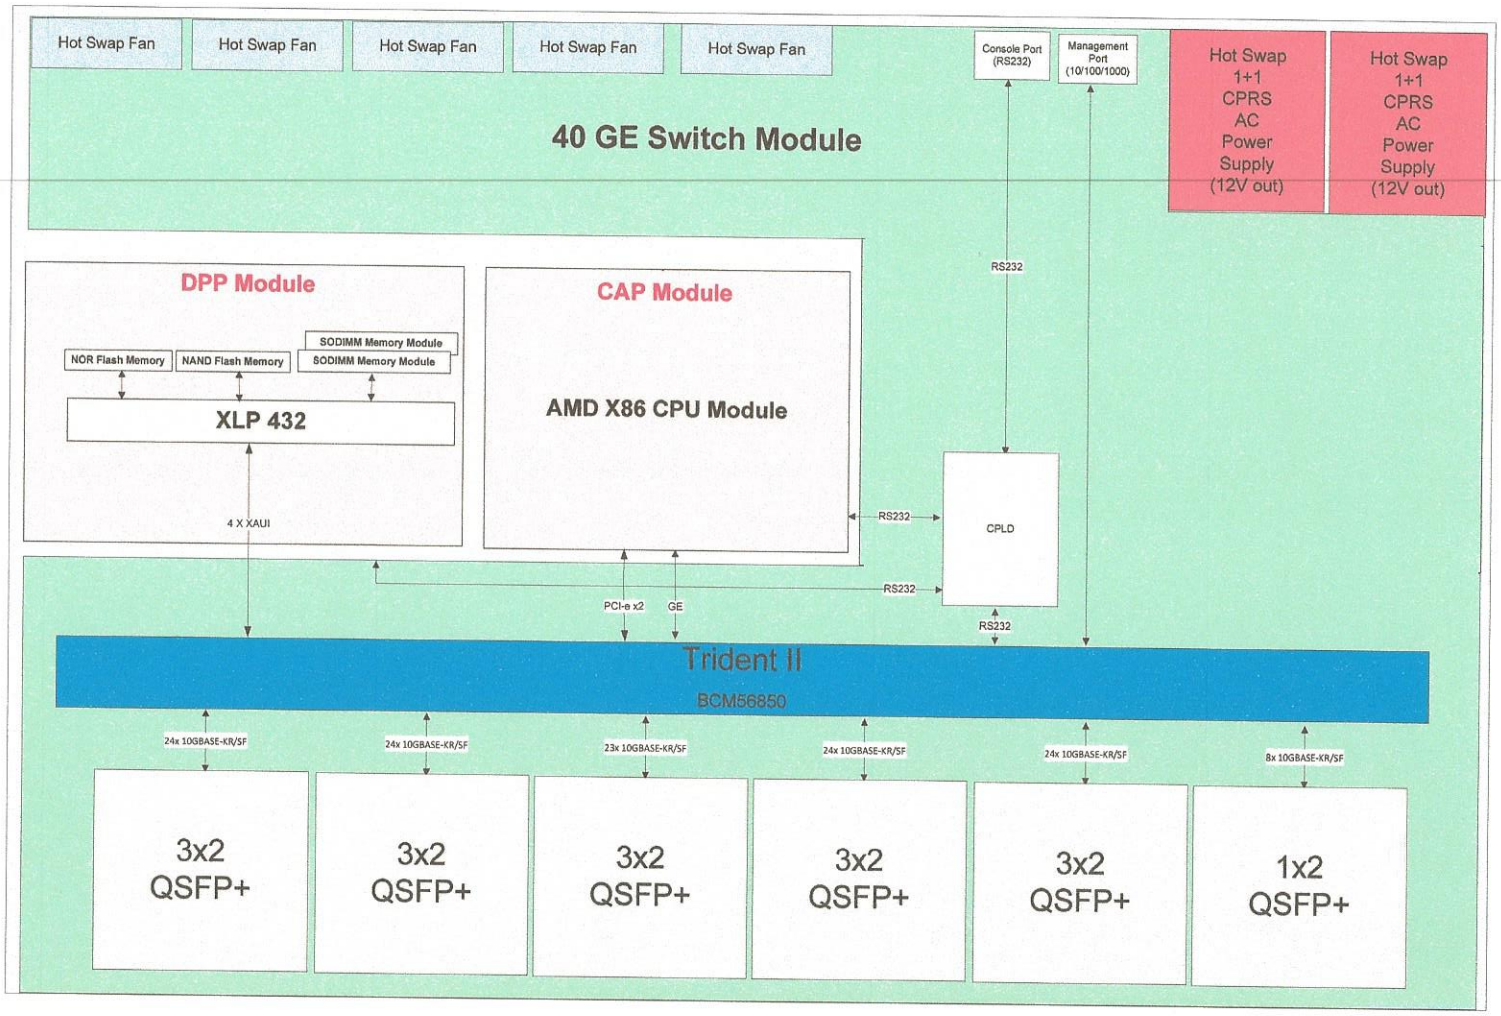

Figure 3 40GbE OCP Open Switch Platform Block Diagram for Standard 19 Inch Rack

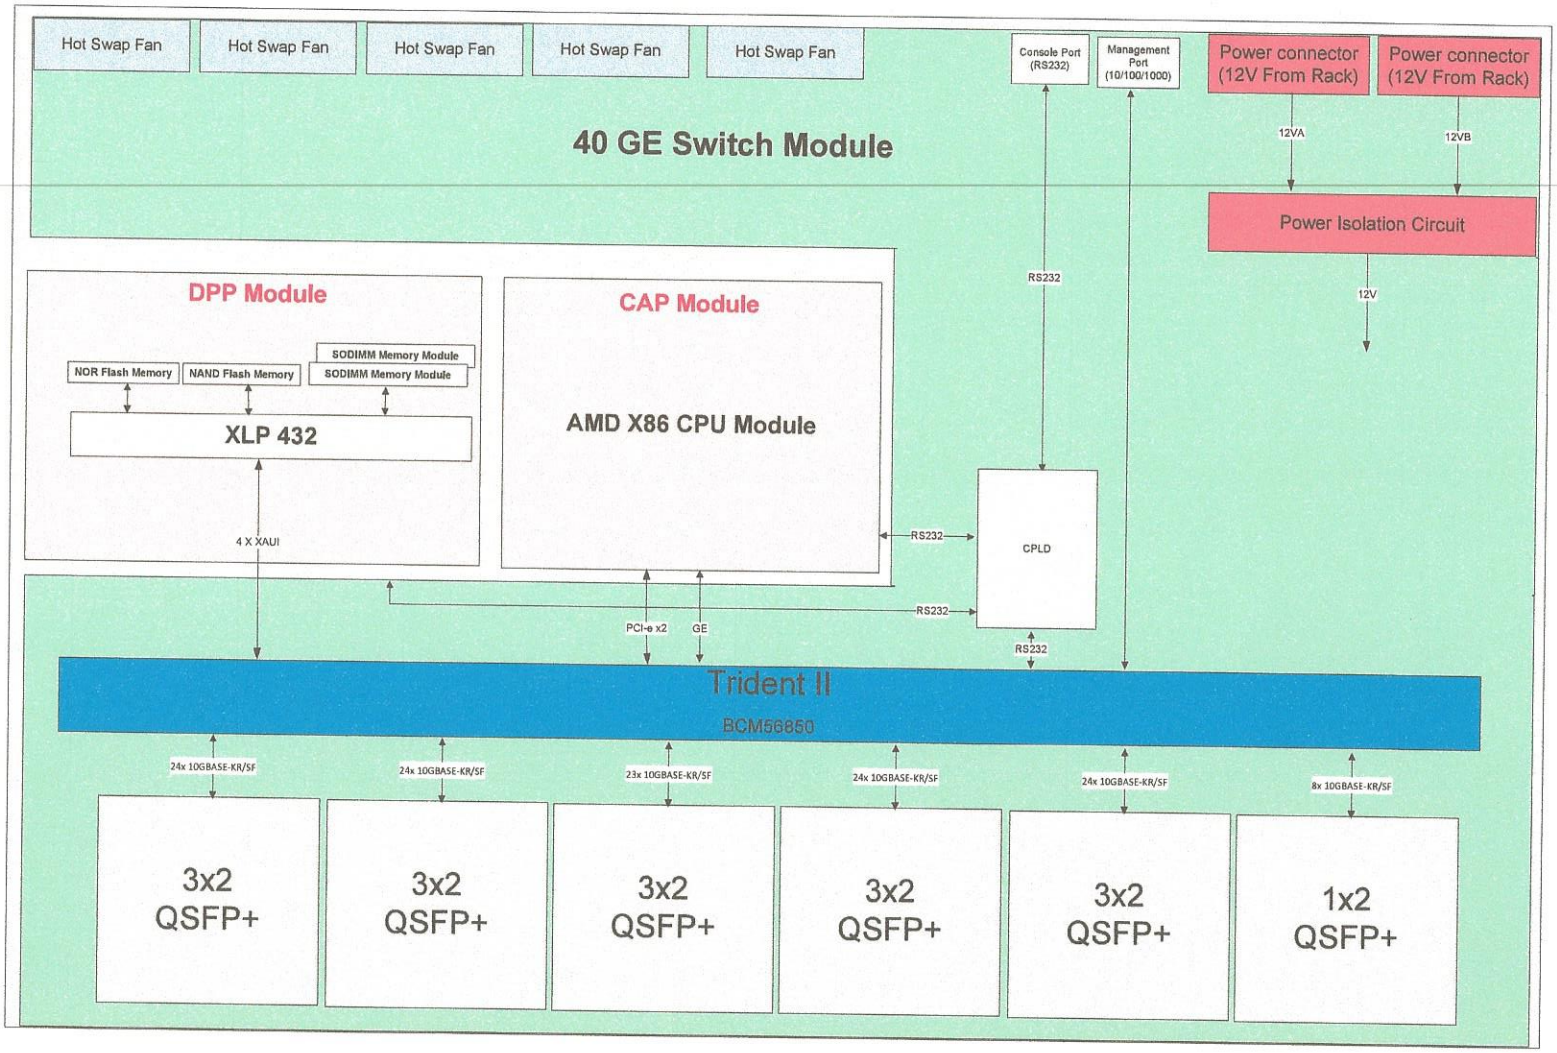

Figure 4 40GbE OCP Open Switch Platform Block Diagram for Open Compute Rack

5 The Trident II Switch Board (TIIS)

The Trident II Switch (TIIS) Board is the main board that interfaces to the CAP, DPP & IO modules. The TIIS Module is based on Broadcom StrataXGS® Trident II switching technology.

The length and width of the Trident II Switch (TIIS) Board in inches and millimeters is shown in Table 2 below.

	Inch	Mm
Length	10.15	257.81
Width	16.7	424.18

Table 2: TIIS Board size

The Trident II Switch (TIIS) Board outline is shown in Figure 5 below.

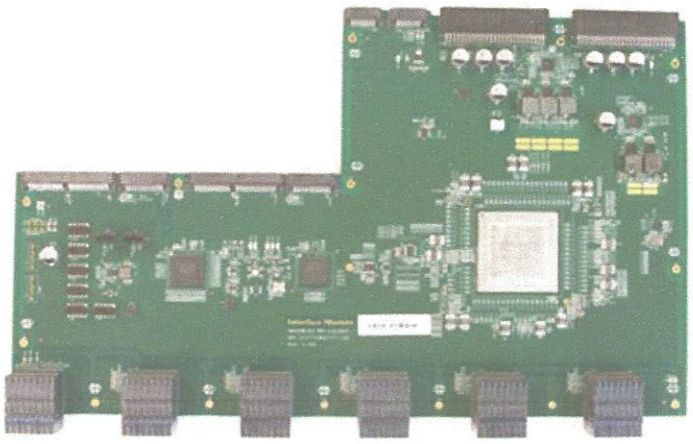

Figure 5: TIIS Switch board outline

The Trident II Switch (TIIS) Board design is based on following major components as shown on Table 3 below.

Item	Description	Manufacturer	Part Number
1	10/40GbE Switch	Broadcom	BCM56850
2	SERDES Repeater	Broadcom	BCM84328
3	CPLD	Altera	EPM2210F324C5N

Table 3: TIIS Major Components

In order to address various combinations of IO modules two models of the TIIS have been design. The 1st option supports combination of 10GbE & 40GbE IO modules ports while the 2nd option only supports the 40GbE IO modules.

5.1 TIIS Option A for Open Leaf Switch

Figure 6 shows a block diagram of the Trident II Switch (TIIS) Board Option A. The BCM56850 Trident II ASIC provides the high performance Ethernet switching functions.

The BCM56850 interfaces to four 10GbE IO Module Connectors for 48 ports of 10GbE connectivity. The connection to each 10GbE IO Module Connector comprises 12x 10G KR for data traffic and MDC/MDIO-based physical layer (PHY) device register read and writes access. Each 10GbE IO Module Connector supports up to 3 Quad SFP+ 10GbE IO Modules.

The BCM56850 interfaces to two 40GbE IO Module Connectors for 12 ports of 40GbE connectivity. The connection to each 40GbE IO Module Connector comprises 6x XLAUI links for data traffic. Physical layer (PHY) device register read and writes access is provided using I2C interfaces connected to the CPLD Module which is connected to the Trident II ASIC using a UART interface. Each 40GbE IO Module Connector supports 1 Hex 40GbE IO Module.

The BCM56850 interfaces with the CAP and DPP Modules using the CAP Connector and DPP Connector respectively.

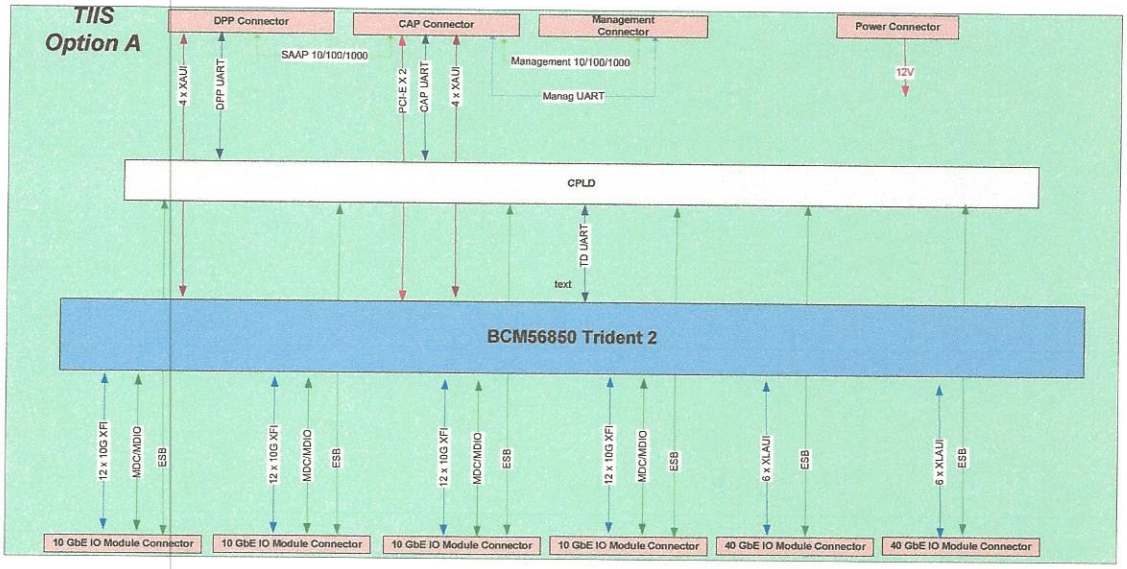

Figure 6: Trident II Switch Module (TIIS) Option A

There are three links to the CAP Module via the CAP Connector which is COMx (CPU On Module) compliant – a PCIe Gen2 x2, a 10/100/1000 Ethernet and a UART connection via the CPLD Module. The PCIe Gen2 interface is used for switch control plane management functionality. Optionally, a 10G XAUI interface can also be supported using the CAP Connector for data traffic between the CPU and the switch. Any COMX compliant CPU Module can be connected to the CAP Connector.

There are two links to the DPP Module via the DPP Connector– a 4x XAUI for up to 40G of data traffic, and a UART connection via the CPLD Module.

In addition, there is a Management Connector for enabling 10/100/1000 Mbps Ethernet based out of band management of the Open Switch Platform. This Management Connector is connected to the Trident II switch using an SGMII and UART links (through the CPLD Module).

5.2 TIIS Option B for Open Spine Switch

Figure 7 shows a block diagram of the Trident II Switch (TIIS) Board Option B. The BCM56850 Trident II ASIC provides the high performance Ethernet switching functions.

The BCM56850 interfaces to six 40GbE IO Module Connectors for 32 ports of 40GbE connectivity. The connection to the 1st five 40GbE IO Module Connector comprises 6x XLAUI links & the 6th 40GbE IO module comprises 2x XLAUI for data traffic. Physical layer (PHY) device register read and writes access is provided using I2C interfaces connected to the CPLD Module which is connected to the Trident II ASIC using a UART interface.

The BCM56850 interfaces with the CAP using the CAP Connector.

Please note that TIIS Option B does not support DPP module.

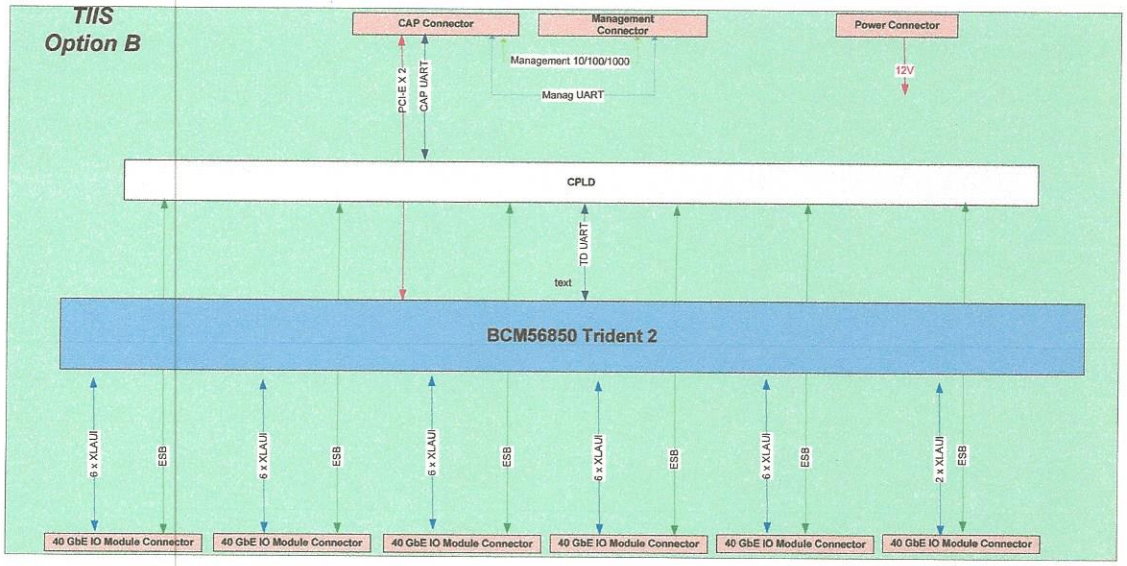

Figure 7: Trident II Switch Module (TIIS) Option B

There are three links to the CAP Module via the CAP Connector which is COMx (CPU On Module) compliant – a PCIe Gen2 x2, a 10/100/1000 Ethernet and a UART connection via the CPLD Module. The PCIe Gen2 interface is used for switch control plane management functionality. Any COMX compliant CPU Module can be connected to the CAP Connector.

In addition, there is a Management Connector for enabling 10/100/1000 Mbps Ethernet based out of band management of the Open Switch Platform. This Management Connector is connected to the Trident II switch using an SGMII and UART links (through the CPLD Module).

6 The Control and Application Processor Module (CAP)

The Open Switch Platform is designed to enable use of exiting CPU Modules available in the market that use the CPU on Module (COMx) interface. There are many suppliers that offer COMx modules with large selection of CPUs for the embedded industry. The CPU Module contains a Control & Application processor for enabling switch control plane functionality. The COMx module is installed into a carrier board that interfaces with the rest of the system.

To enable standard and server class Linux based operating systems in the Open Switch Platform, the default Control & Application processor module is based on AMD G Series X86 CPU. Both dual core and quad core options are supported. AMD's G-Series System on a Chip (SOC) combines a low-power CPU and I/O controller onto a single chip, delivering performance and efficiency without feature compromise. Besides, it supports standard and server class Linux based operating systems such as Ubuntu or CentOS.

The CAP module runs Broadcom Switch SDK software and network operating systems to program the Trident II switch and also provide a vehicle to execute customer applications.

The CAP module resides on a carrier board and interfaces to the Trident II Switch (TIIS) Board.

The length and width of the CAP module in inches and millimeters is shown in Table 4 below.

	Inch	mm
Length	4.92	125.00
Width	3.74	95.00

Table 4: CAP Board size

The CAP module outline is shown in Figure 8 below.

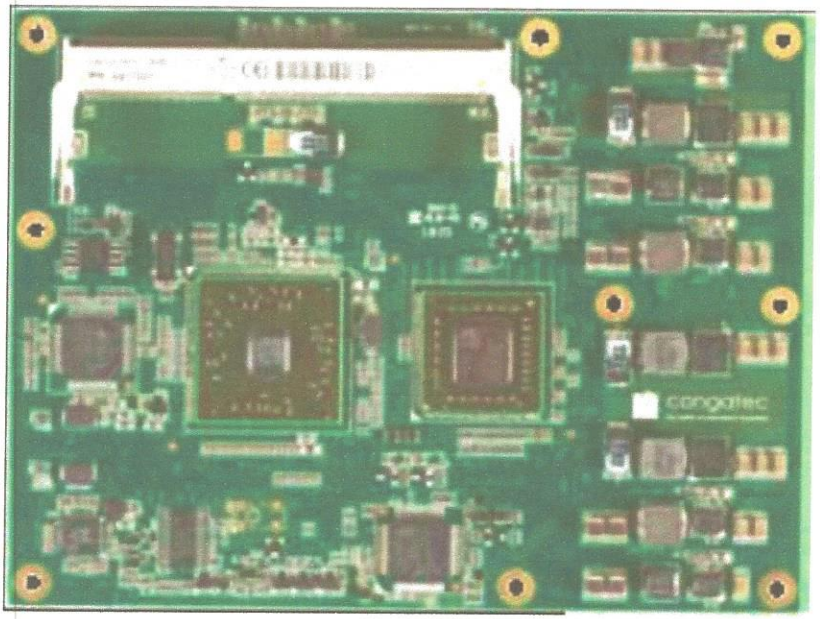

Figure 8: CAP Board outline

The CAP module design is based on the following major components as shown in Table 5.

Item	Description	Manufacturer	Part Number
1	X86 CPU	AMD	T40N
2	IO Hub	AMD	A55E FCH
3	SODIMM Memory	Micron Technology	MT18KSF25672HZ
4	GE MAC/PHY Controller	Broadcom	BCM5725

Table 5: CAP Module Major Components

The AMD G Series x86 based CAP Module in the COMx form factor is shown in Figure 6. The CAP Module exposes four interfaces – a PCIe Gen2 x2 for control plane connectivity to the Trident II switch ASIC, a GE interface for management traffic, a UARTx2 interface and also a mini SATA interface additional storage requirements. The PCIe and UART interfaces are provided using the IO HUB A55E FCH chip that acts as a companion chip to the AMD x86 CPU, connected via an UMIx4 interface. The GE connectivity is provided using a Broadcom GE MAC/PHY Controller BCM84328 solution that connects to the AMD x86 processor using PCIe interface. The GE connection provides flexible baseboard management capabilities.

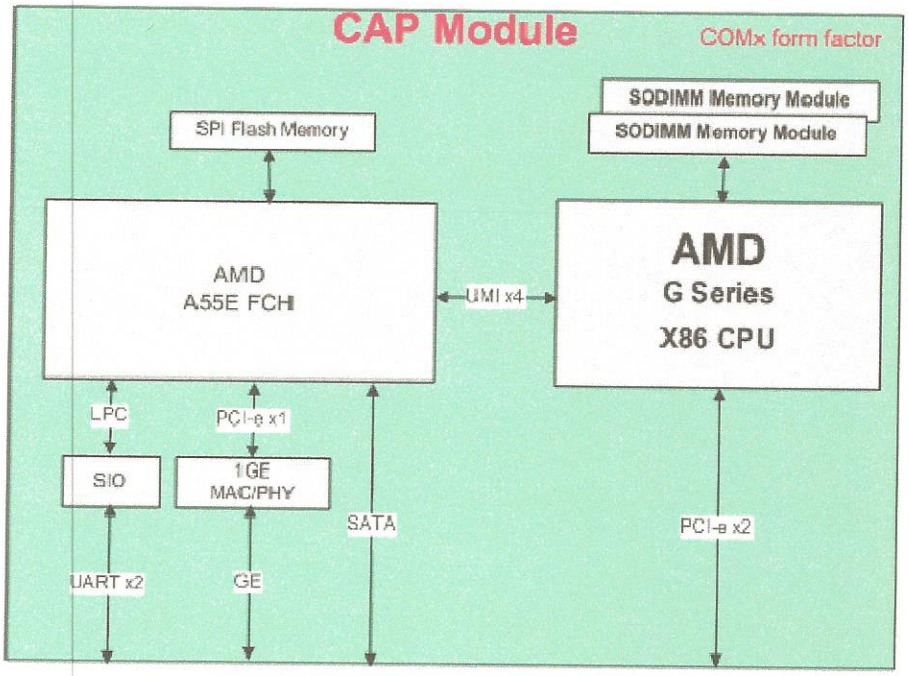

Figure 9: Control and Application Processor Module (CAP)

Flash and Memory connectivity is provided using the A55E FCH IO HUB and G Series x86 CPU chips. Flash Memory size of up to 512 MB is supported using the SPI and mini SATA interfaces available on the IO HUB chip. Minimum Flash Memory Size supported is 64 MB. SDRAM based memory is supported using the Micron SODIMM Memory Modules connected to the AMD G Series CPU. 2GB of SDRAM is supported on this CAP Module.

7 The Data Plane Processor Module (DPP)

In the Open Switch Platform, the optional Data Plane Processor (DPP) Module is provided to enable specialized network functions such as L4-L7 processing. The DPP module is based on Broadcom XLP432 CPU. The Broadcom XLP400 Series Processors, which includes the XLP432 CPU, are highly-scalable devices that incorporate key functions of a high-end communication system, including wired and wireless security, networking, storage, data center acceleration, load balancing, and other acceleration engines.

The DPP Module board size is shown in Table 6 below.

	Inch	mm
Length	6.75	171.45
Width	4.00	101.6

Table 6: DPP Module Board Size

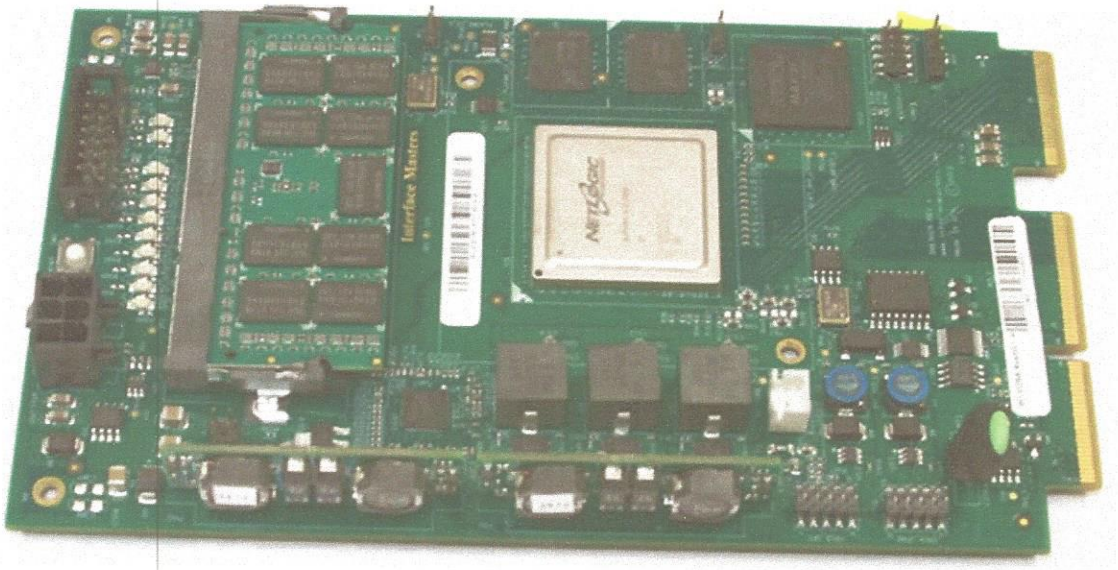

Figure 10 DPP Module Outline

The DPP Module design is based on following major components:

Item	Description	Manufacturer	Part Number
1	32 Core MIPS CPU	Broadcom	XLP432
2	SODIMM Memory	Micron Technology	MT18KSF25672HZ
3	CPLD	Altera	EPM2210F324C5N

Table 7: DPP Module Major Components

As shown in Figure 11, the DPP Module exposes the following interfaces for IO and storage:

- 4x XAUI interfaces for up to 40G of data traffic connectivity. The high throughput connectivity can be used to provide data plane co-processing functions in conjunction with the Trident II switch chip
- A UART interface
- Interfaces for Flash and Memory subsystems. Up to 16GB of DDR3 RAM and 2GB of Flash Memory are supported.

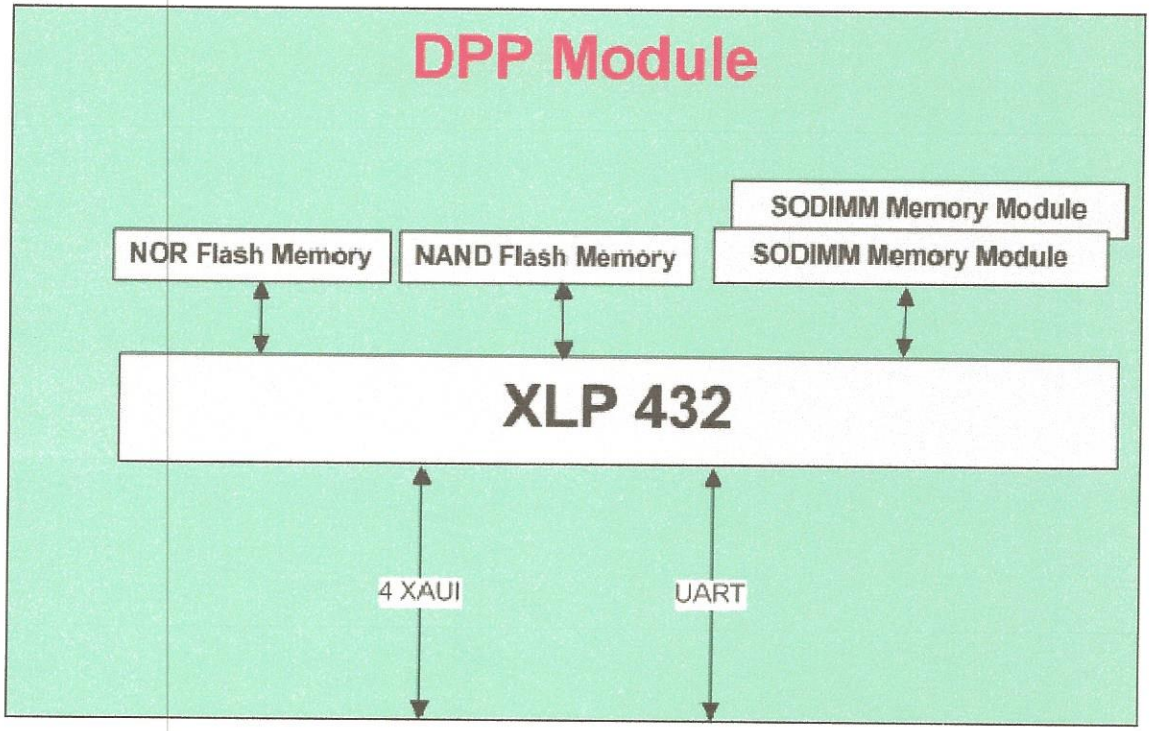

Figure 11: Data Plane Processor Module (DPP)

8 Quad 10GbE IO Module

The Open Switch Platform supports up to twelve 10GbE IO Modules. Each IO module supports up to four dual speed SFP+ modules. The 10GbE IO Module board size is shown in Table 8. The board outline is shown in Figure 12.

	Inch	mm
Length	6.75	171.45
Width	4.00	101.6

Table 8: 10GbE IO Module board size

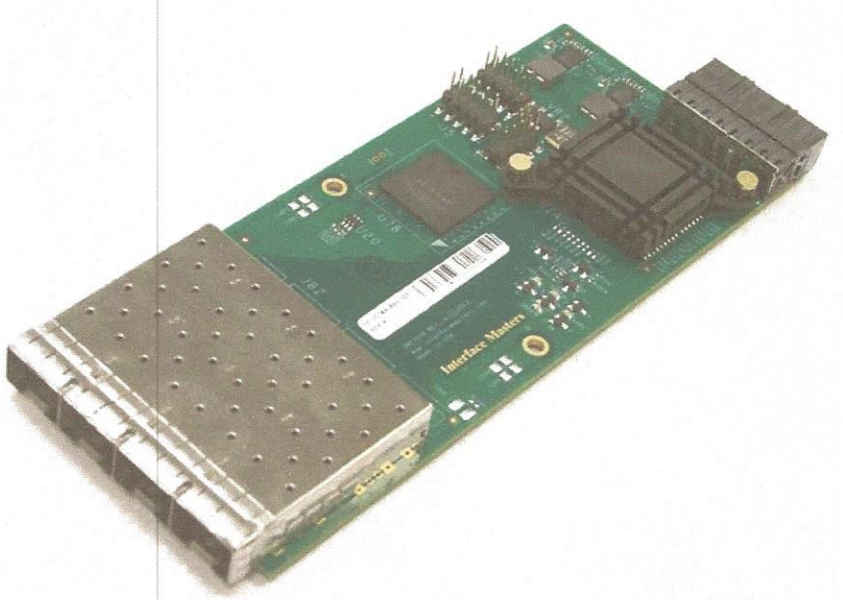

Figure 12: 10G IO board outline

The 10GbE IO Module design is based on following major components:

Item	Description	Manufacturer	Part Number
1	SERDES Repeater	Broadcom	BCM84328

Table 9: 10GbE IO Module Major Components

As shown in Figure 13, the 10GbE IO Module exposes the following interfaces:

- 4 SFP+ interfaces for external facing 10GbE connectivity
- 4 10G KR interfaces that connect to the High Speed Connector. High Speed Connector connects to the Trident II switch ASIC in the main board.

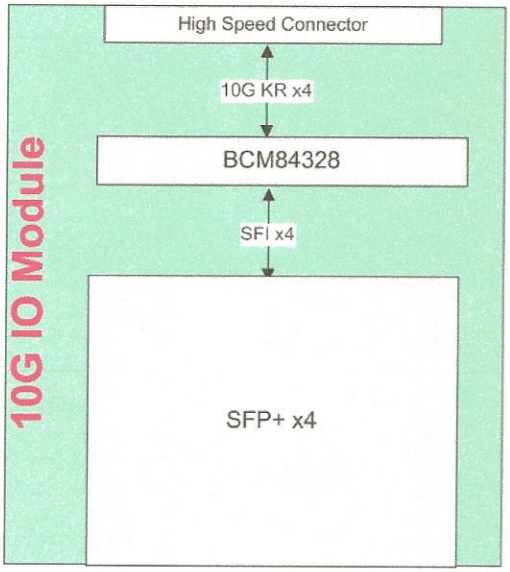

Figure 13: Quad 10GbE IO Module

9 Hex 40GbE IO Module

The Open Switch Platform supports up to two 40GbE IO Modules. Each IO module supports up to six QSFP+ modules. The 40GbE IO Module board size is shown in Table 10. The board outline is shown in Figure 14.

	Inch	mm
Length	6.75	171.45
Width	4.00	101.6

Table 10: 40G IO board size

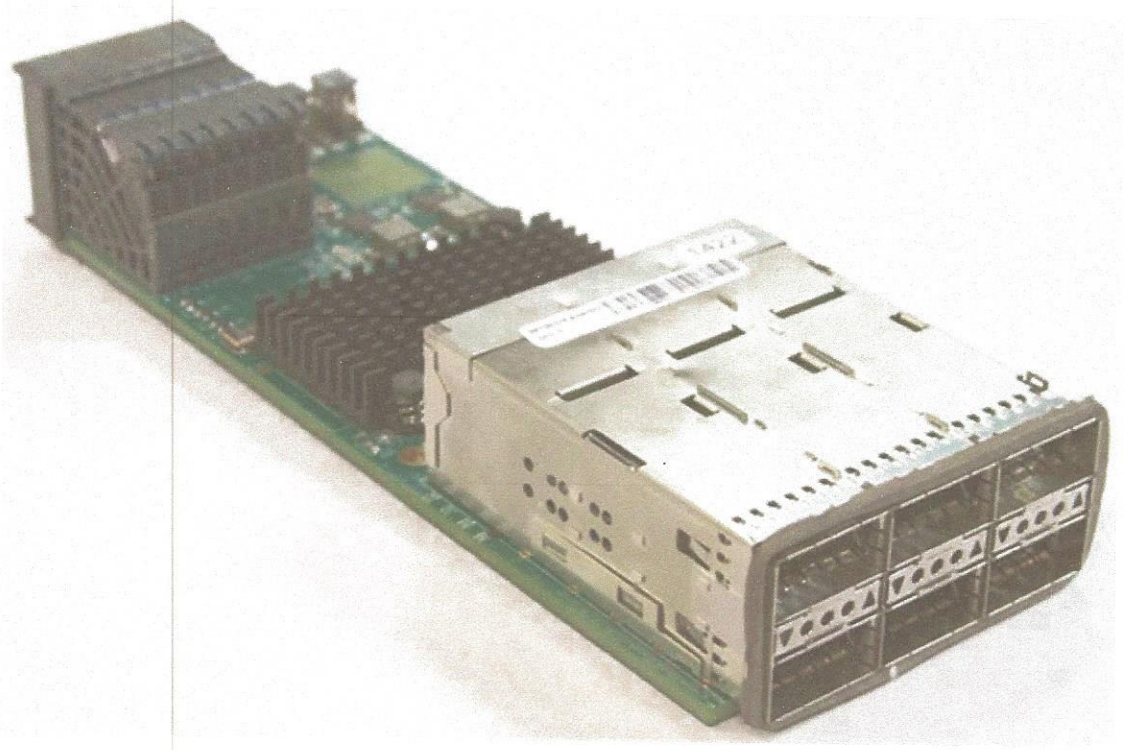

Figure 14 40G IO board outline

The Open Switch Platform 40GbE IO Module design is based on the following major components:

Item	Description	Manufacturer	Part Number
1	SERDES Repeater	Broadcom	BCM84328

Table 11: 40G IO Major Components

As shown in Figure 15, the 40GbE IO Module exposes the following interfaces using 3 BCM84328 SERDES Repeater chips:

- Each BCM84328 exposes 2 QSFP+ interfaces using 2x XLPPI, for external facing 40GbE connectivity. Together, the three BCM84328 exposes 6 QSFP+ interfaces
- Each BCM84328 exposes 2x XLAUI to the High Speed Connector. Together, the three Each BCM84328 exposes 6x XLAUI interfaces to the High Speed Connector that connects to the Trident II switch ASIC in the main board.

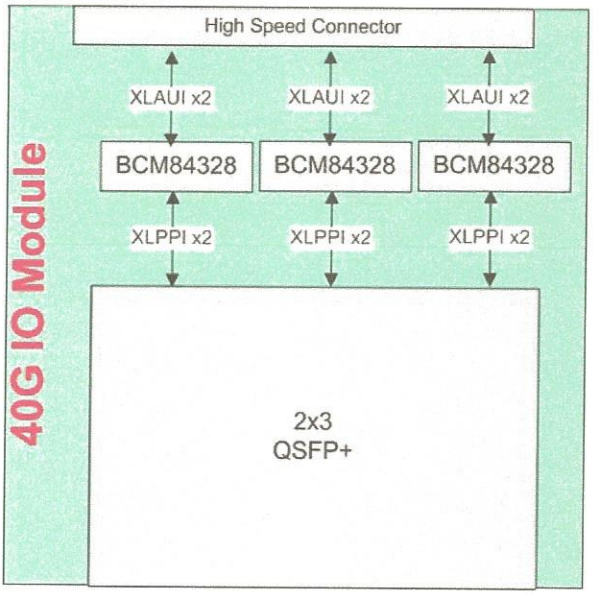

Figure 15: Hex 40G IO Module

10 Leveraged Design

This OCP Open Switch Specification and Platform have been designed to be easily leveraged to support many data center and service provider networking switch applications, both as a Top of Rack (ToR) Switch or as an Aggregation Switch. In CLOS or flat network topologies, the Open Switch Specification applies to both Spine and Leaf switches. Specifically, use cases such as SDN Switch, General Purpose High Density ToR/Leaf Switch, General Purpose High Density Spine Switch and Specialized Network Appliance Switch can be supported. All of the four use cases can be built around the same core design, but a few changes in both component/module population and the PCB can be made to optimize the design for the given application. The Trident II Switch (TIIS) module forms the base module around which the CAP and DPP modules can be populated or customized to achieve application specific optimizations. The below table shows use case and design optimization examples.

Application	Use Case	Component Change
SDN Switch	Enable extensibility of the Switch network OS by providing a Server Linux class development environment. Data Center and Service Provider Operators can use the development environment to add custom applications to run on the switch platform to enhance network functions such as network automation, instrumentation and traffic engineering.	Uses the TIIS and AMD x86 CPU based CAP modules only. The AMD x86 processor, BSP and Broadcom switch SDK software running on Ubuntu Linux provides the necessary development environment. Multiple network OS and protocol options can be supported.
OpenFlow 1.3 Switch	Using OpenFlow 1.3+ software abstraction layers provided with the Broadcom switch SDK, a scalable multi-table implementation can be	Uses the TIIS and AMD x86 or BRCM XLP316 CPU based CAP modules. Processor specific BSP and Broadcom

	enabled to support popular OpenFlow 1.3+ use cases such as network virtualization, network utilization optimization, dynamic reconfiguration, network visibility and WAN integration	switch SDK software running on popular embedded Linux provides the base platform over with OpenFlow Agents and other network OS and protocol options can be installed
General Purpose High Density ToR/Leaf Switch	Enable highly scalable flat L2 or L3 networking topologies with multipathing. Use high radix switches in spine/leaf topologies to build cost-effective CLOS-based topologies. Leverage IP/L3 based networking for building Internet scale Web 2.0 topologies using high L3 and ECMP scale.	Uses the TIIS and CAP modules. An alternate FreeScale PPC or BRCM XLP MIPS based CAP Module may be used instead of the x86 CPU for higher performance or better compatibility with existing control plane software from different vendors. The number of 40GbE QSFP connectors may be depopulated to support six ports of 40GbE only to serve as uplinks. QSFP to SFP+ breakout cables are recommended to enable 24 10GbE uplink ports.
General Purpose High Density Aggregation/Spine Switch	Same as ToR/Leaf Switch	Uses the TIIS Option B and CAP modules. An alternate FreeScale PPC or BRCM XLP MIPS based CAP Module may be used instead of the x86 CPU for higher performance or better compatibility with existing control plane software from different vendors. The design supports up to 32 ports of 40GbE QSFP that

		can be connected to the 40GbE QSFP uplinks of a Leaf Switch. The number of 40GbE QSFP connectors may be depopulated to support fewer 40GbE ports if needed. QSFP to SFP+ breakout cables can be used to enable a mix of 10 and 40GbE ports. For example, configurations such as 96 ports 10GbE and 8 ports of 40GbE can be created.
Specialized Network Appliance Switch	The Trident II switch provides scalable L2-L3 data plane processing. The XLP432 CPU as a data plane co-processor extends the processing capabilities to L4-L7, enabling applications such as Firewall, Load Balancer, Deep Packet Inspection for Data Center network infrastructures, and Mobility Management and Gateways for Evolved Packet Core network infrastructures.	Uses the TIIS, CAP and DPP modules. The 4xXAUI links between the TIIS and DPP are used to high performance packet I/O for split data plane processing between the Trident II switch and the XLP432 CPU.

11 Chassis Information

The TOR/Leaf and the Spine switch share the same physical chassis but with different front panel. The front panel for TOR switch accommodates a combination of 10G & 40G IO modules while the front panel for the spines switch supports only the 40G IO modules.

Figure 16: OCP Open Switch Platform Chassis for Leaf Switch - Front View

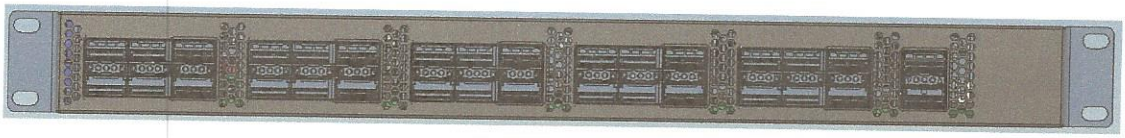

Figure 17: OCP Open Switch Platform Chassis for Spine Switch - Front View

11.1 Physical Dimension

Side	Inch	MM
Width	17.25	438.15
Depth	24.5	622.30
Height	1.72	43.688

Table 12: Physical Dimension

11.2 Top of the Rack / Leaf Configuration

Figure 18 below shows an implementation (by Interface Masters Technologies, Inc.) of the 19 inch rack version of the Open Switch Platform, with the cover removed to show how the components/modules of the platform are laid out.

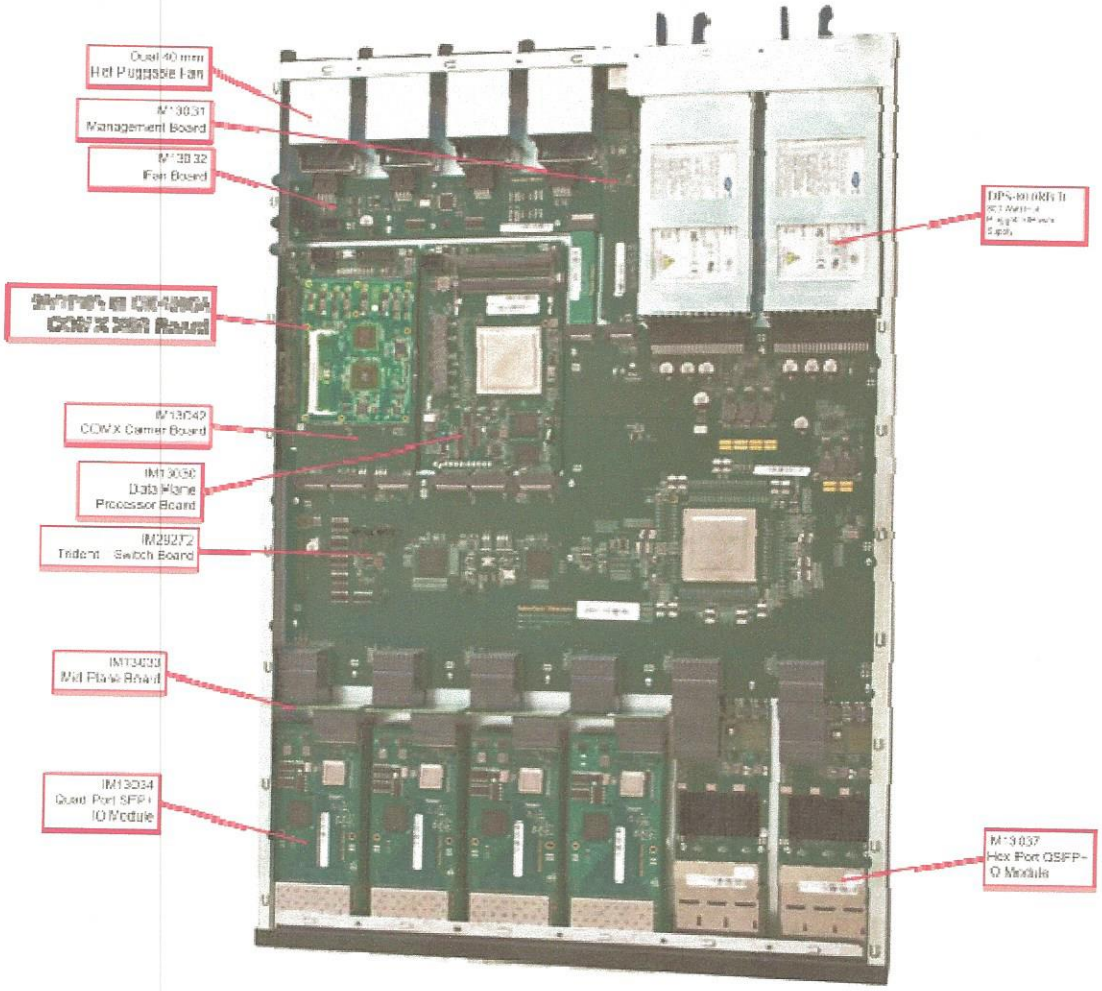

Figure 18: Physical Layout 19 inch rack version of the Open Leaf Switch Platform

Figure 19 below shows an implementation (by Interface Masters Technologies, Inc.) of the Open Switch Specification in a Chassis that is 1U high and 26 inched deep. The front panel shows a fully populated 10 and 40GbE IO Modules configuration, exposing 48 ports of 10GbE and 12 ports of 40GbE

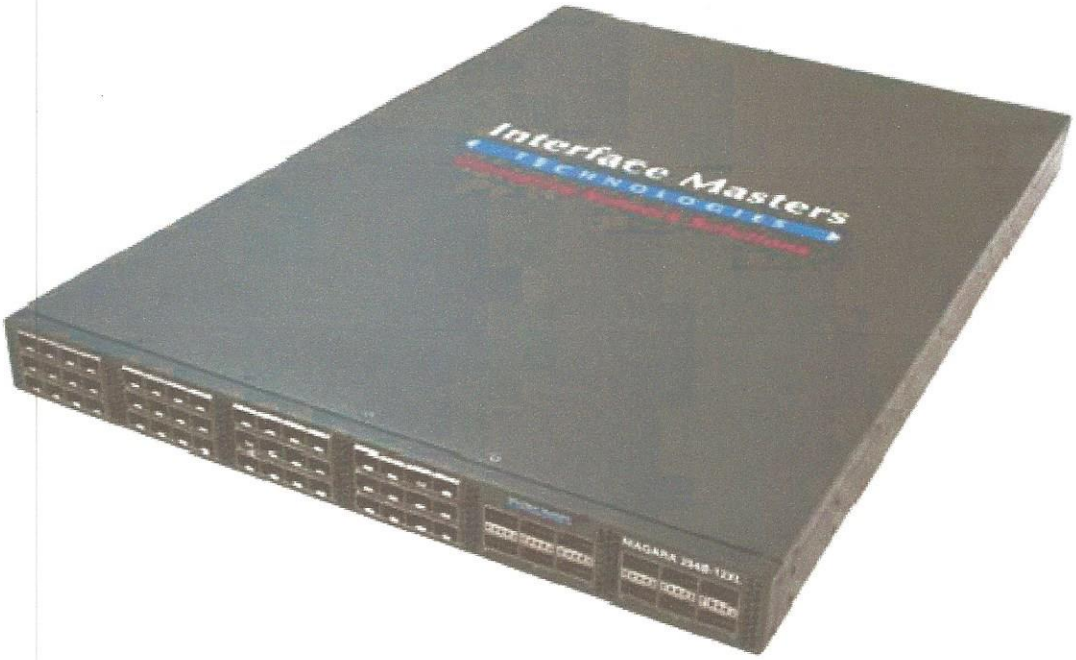

Figure 19: Example of OCP Open Switch Chassis for a Leaf Switch

Figures below show close ups of the front, rear and sides of the chassis for the Leaf Switch design:

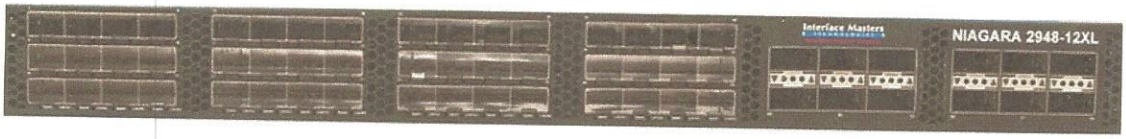

Figure 20: OCP Open Switch Platform Chassis for Leaf Switch - Front View

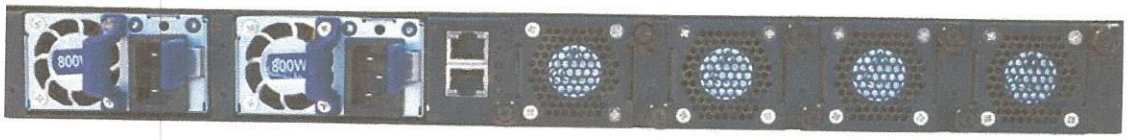

Figure 21: OCP Open Switch Platform Chassis for Leaf Switch - Rear View

Figure 22: OCP Open Switch Platform Chassis for Leaf Switch - Side View

11.3 Spine Configuration

Figure 23 below shows an implementation (by Interface Masters Technologies, Inc.) of the 19 inch rack version of the Open Switch Platform for a spine Switch with TIIS Option B, with the cover removed to show how the components/modules of the platform are laid out.

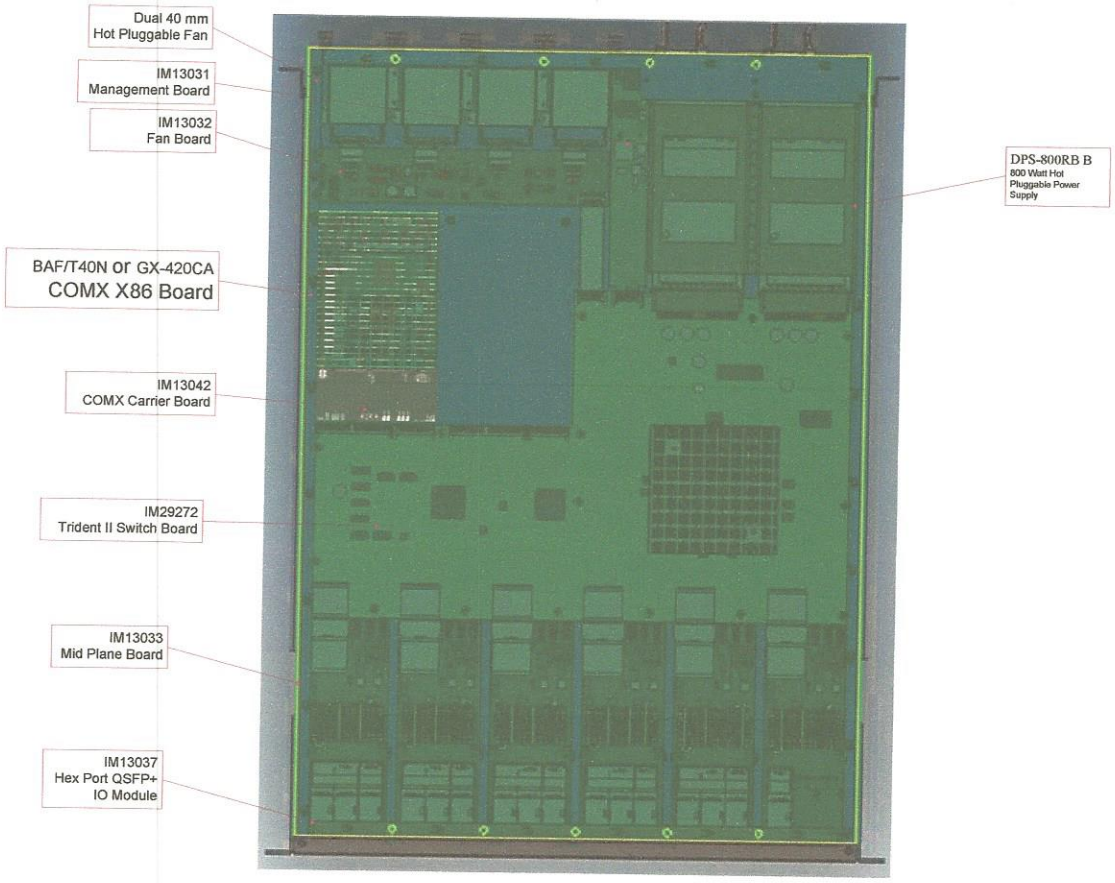

Figure 23: Physical Layout 19 inch rack version of the Open Switch Platform for a Spine Switch

Figure 24 below shows an implementation (by Interface Masters Technologies, Inc.) of the Open Switch Specification for a spine Switch in a Chassis that is 1U high and 26 inched deep. The front panel shows a fully populated 40GbE IO Modules configuration, exposing 32 ports of 40GbE.

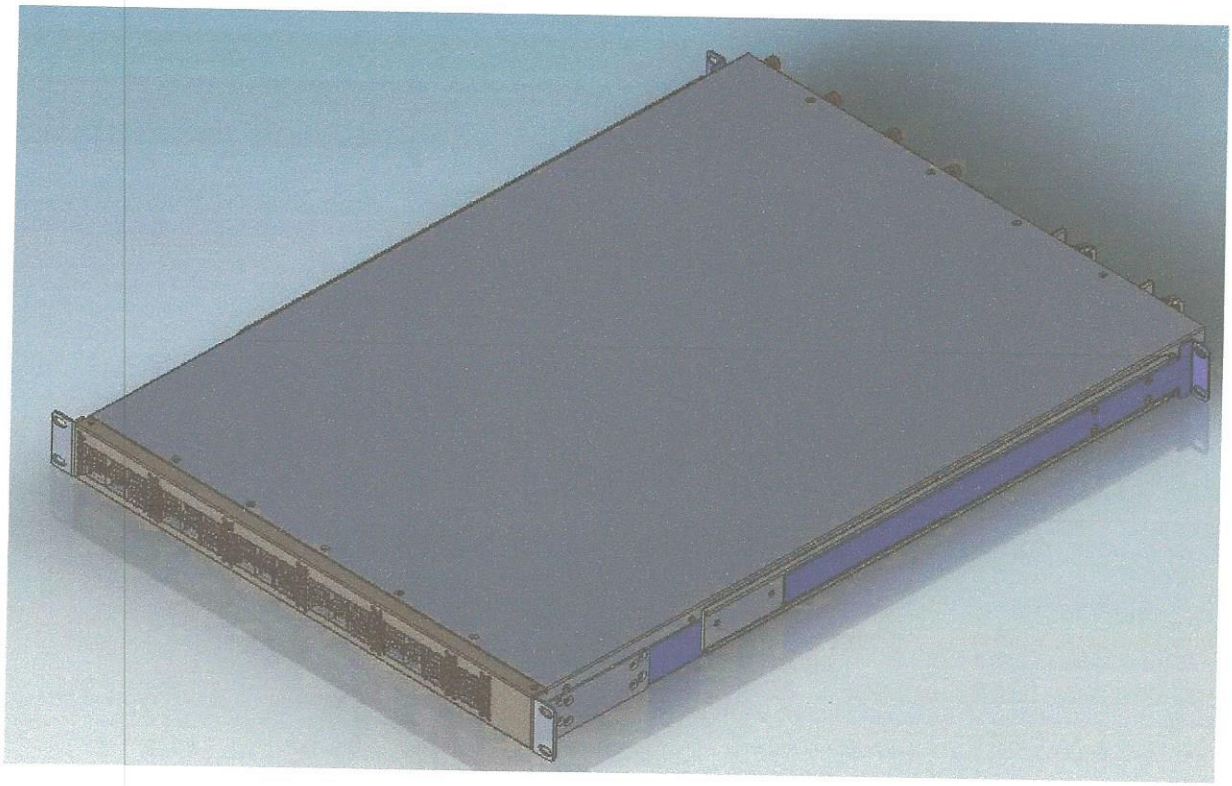

Figure 24: Example of OCP Open Switch Chassis for a Spine Switch

Figures below show close ups of the front, rear and sides of the chassis for the Spine Switch design:

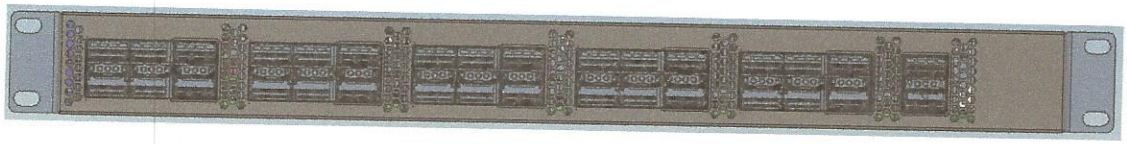

Figure 25: OCP Open Switch Platform Chassis for Spine Switch - Front View

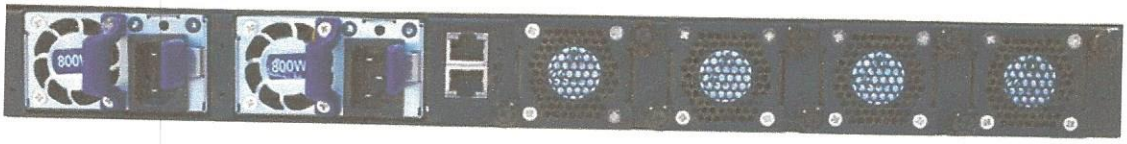

Figure 26: OCP Open Switch Platform Chassis for Spine Switch - Rear View

Figure 27: OCP Open Switch Platform Chassis for Spine Switch - Side View

12 CAP AMD CPU BIOS

PXE Boot

The BIOS for the AMD x86 G-Series CPU supports PXE boot. When PXE booting, the system first attempts to PXE boot from the first Ethernet interface (eth0). If a PXE boot on the first Ethernet interface fails, the BIOS attempts to PXE boot from the second Ethernet interface (eth1).

Other Boot Options

The BIOS also supports booting from SATA/SAS and USB interfaces. The BIOS provides the capability to select boot options.

- Can be scripted and propagated to multiple machines

13 System Management

13.1 Ethernet Port

One 10/100/1000 GE port is provided to remotely setup and control the system.

This port is directly connected to the Control Application Processor.

13.2 Console Port

One RS232 port is provided for local setup and control the system.

This port is directly connected to the Control Application Processor.

14 Fan

There are four sets of dual stacked hot swappable fan in the system.

The temperature inside the system is monitored by a set of temperature sensors which is received by the Control Application Processor, CAP.

The CAP controls the speed of the fans depending on the temperature sensors and the system configuration settings.

15 Power System

15.1 Power Supply for 19 inch Rack

The 19 inch rack version of the Leaf & Spine Switch system deploys two 600 Watt hot swappable 1+1 redundant power supplies. The power supplies are part for CRPS series which are sourced from Sparkle, Delta and several other power supply manufacturers. These CRPS power supplies offer 400W to 800W within the same physical dimension and interface. These Power supplies are 80 plus gold efficient and support PMBus specification 2.0

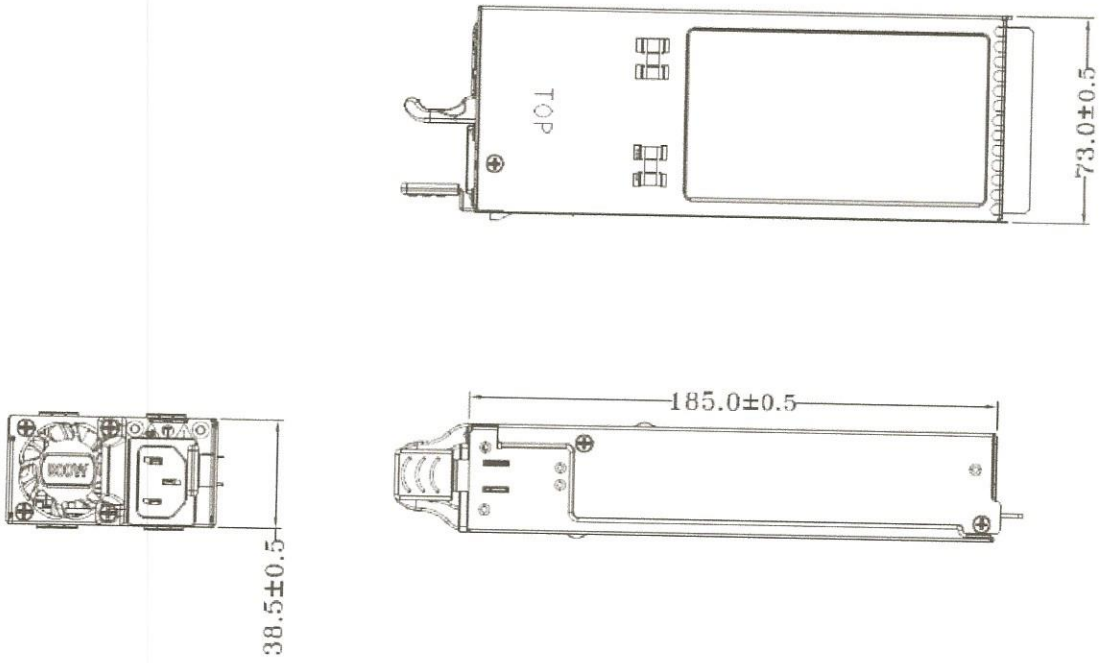

• Figure 28 Power Supply

15.2 Support for Open Compute Infrastructure

The system will support the Open compute power infrastructure by directly receiving +12V from the Open Rack.

15.3 Power Budget

The power budget for the system depends greatly on following:

a) Speed of the Control Application Processor, CAP
b) Speed and memory size of the Data Plane Processor, DPP
c) Number of Active 10 GE and 40 GE ports

16 Environmental Requirements

The Open Switch Platform meets the following environmental requirements:

- Gaseous Contamination: Severity Level G1 per ANSI/ISA 71.04-1985
- Ambient operating temperature range: -5°C to +45°C.
- Operating and storage relative humidity: 10% to 90% (non-condensing)
- Storage temperature range: -40°C to +70°.
- Transportation temperature range: -55°C to +85°C (short-term storage).

In addition, the full system has an operating altitude with no de-ratings of 1000m (3300 feet).

17 Regulatory Compliance

The Open Switch Platform meets the following regulatory compliances:

- FCC/CE Class "A" EMC
- UL Safety Enabled

17.1 Vibration and Shock

The Open Switch Platform meets shock and vibration requirements according to the following IEC specifications: IEC78-2-(*) and IEC721-3-(*) Standard & Levels. The testing requirements are listed below.

Operating Non-Operating

Vibration 0.5g acceleration, 1.5mm amplitude, 5 to 500 Hz, 10 sweeps at 1 octave/minute for each of the three axes (one sweep is 5 to 500 to 5 Hz) 1g acceleration, 3mm amplitude, 5 to 500 Hz, 10 sweeps at 1 octave/minute for each of the three axes (one sweep is 5 to 500 to 5 Hz)

Shock 6g, half-sine 11mS, 5 shocks for each of the three axes 12g, half-sine 11mS, 10 shocks for each of the three axes

BROADCOM CORPORATION

By: _____

Name: SHRIRAM VELAGA

Title: VP/GM, Network Switch

Date: 11/08/13