



OPEN
Compute Project

Barefoot Networks

Wedge 100B Series of Switches

2RU 65x100GE switch

&

1RU 32x100GE switch

Author: Barefoot Networks



© 2017 Barefoot Networks, Inc.

As of July 26, 2016, the following persons or entities have made this Specification available under the Open Compute Project Hardware License (Permissive) Version 1.0 (OCPHL-P), which is available at <http://www.opencompute.org/community/get-involved/spec-submission-process/>.

Barefoot Networks, Inc.

Your use of this Specification may be subject to other third party rights. THIS SPECIFICATION IS PROVIDED "AS IS." The contributors expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to the Specification. The Specification implementer and user assume the entire risk as to implementing or otherwise using the Specification. IN NO EVENT WILL ANY PARTY BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS SPECIFICATION OR ITS GOVERNING AGREEMENT, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND WHETHER OR NOT THE OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Document Revision History

Revision	Date	Description
0.81	1/22/2017	Initial Release
0.82	2/10/2017	Minor updates



Contents

1	Scope	5
2	Overview	6
3	Feature Highlights	7
4	Physical Overview	8
5	Field Replaceable Components	11
5.1	Power Supply Unit	11
5.2	Fan Modules	13
6	System Overview	14
6.1	Main Board	14
6.2	CPU Module	15

1 Scope

This document outlines the technical specifications for the Barefoot Networks' Wedge 100B series of switches submitted to the Open Compute Foundation.



2 Overview

This document describes the technical specifications of the Wedge 100B series of switches designed by Barefoot Networks. The Wedge 100B designs are cost optimized switch designs focused on deployments which support 10G/25G/40G/50G/100G server connectivity and providing 100Gb uplinks.

The Wedge 100BF-32X switch supports thirty-two QSFP28 ports that each can operate at 4x10G or 4x25G with QSFP28-4xSFP28 break out cables, 2x50G with QSFP28-2xQSFP28 break out cables, 40G with standard QSFP+ optics/DAC cables, and 100G with QSFP28 optics/DAC cables.

The Wedge 100BF-65X switch supports sixty-five QSFP28 ports that each can operate at 4x10G or 4x25G with QSFP28-4xSFP28 break out cables, 2x50G with QSFP28-2xQSFP28 break out cables, 40G with standard QSFP+ optics/DAC cables, and 100G with QSFP28 optics/DAC cables.

The Wedge 100B switches use Barefoot Networks' P4 programmable Tofino Ethernet switch ASICs.

3 Feature Highlights

Barefoot Networks' Wedge 100B series of switch platforms consists of 6.5Tbps and 3.2Tbps systems for programmable networking applications.

Both switch platforms are equipped with:

- Intel Xeon D-1517 processor connected via COM Express CPU connector
- PCIe Gen3 x4 connectivity to Tofino
- Board Monitor Controller accessible via front panel RJ-45 or USB connectors
- Redundant (1+1) hot-pluggable AC/DC modules
- Hot-pluggable (N+1) fan modules

Production specifications for Wedge 100B switch platforms are summarized in the table below.

Table 1: Wedge 100B Switch Series Product Specification

Specification	Wedge 100BF-65X	Wedge 100BF-32X
Barefoot Networks Tofino Device	BFN-T10-064Q	BFN-T10-032D
System Throughput	Up to 6.5 Tb/s	Up to 3.2 Tb/s
zQSFP+ Ports	65	32
Maximum 40/100GbE Port Density	65	32
Maximum 50GbE Port Density	130	64
Maximum 10/25GbE Port Density	260	128
Maximum 1GbE Port Density ¹	4	0
CPU	Intel Xeon D-1517, 4 Core@2.2GHz 8GB DDR4 SDRAM, 128GB M.2 SSD	Intel Xeon D-1517, 4 Core@2.2GHz 8GB DDR4 SDRAM, 128GB M.2 SSD
Board Monitor Controller	Aspeed AST1250A1-GP 400MHz	
Front Panel Management Interfaces	RJ-45 1GbE Port to CPU and BMC × 1 RJ-45 UART Port to CPU × 1 USB 2.0 Port to CPU and BMC × 1	
Facebook Debug Interfaces	Front Panel 14-pin debug connector Rear Panel 2×2 RJ45 Ports	
AC Power	Redundant (1+1) hot-pluggable 1100W 110-240VAC AC/DC power supplies	Redundant (1+1) hot-pluggable 850W 110-240VAC AC/DC power supplies
DC Power (OpenRack Power Bus)	N/A	OCP V2 Rack 12V Power Pass-through Card




4 Physical Overview

Table 2: System Physical Specification

Specification	Wedge100BF-65X	Wedge100BF-32X
Rack Units	2 RU	1 RU
Dimensions (WxDxH)	17.32" × 19.97" × 3.482" (440 × 507 × 88.45 mm)	17.32" × 19.97" × 1.732" (440 × 507 × 44 mm)
Weight	~45 lbs	~22 lbs
Environment Specification		
Airflow Direction	Front-to-Back	
Temperature, Operating	0 to 40°C (32 to 104°F)	
Temperature, Storage	-40 to 70°C (-40 to 158°F)	
Relative Humidity, Operating	5 to 90% noncondensing	
Relative Humidity, Storage	5 to 95% noncondensing	
Altitude, Operating	0 to 10,000 ft (0 to 3048m)	

The following figures illustrate the front/back panel interface and board components for Wedge 100B series of switches.






Figure 2: Wedge100BF-65X Component and Interface Overview

5 Field Replaceable Components

5.1 Power Supply Unit

Wedge100BF-65X systems support two hot swappable PSUs plugged in at the same time for redundancy. It requires 1100W AC/DC supplies with internal 40 mm fan unit. The Power supply is PFE1100 from Bel Power Solutions. This is the same form factor Power Supply as Wedge100BF-32X design which uses PFE850 for 850W.

Power supply features.

- Input AC Voltage: 90V-264V AC
- Output DC Voltage: 12V and 5V Standby (5V Standby is not used in the design)
- Hot Swappable
- Load Sharing between the 2 PSU
- Output Connector: Tyco P/N 2-1926736-3
- PCB Mating Connector: Tyco P/N 2-1926733-5
- SMBUS interface for management and control
- Dimension: 2.145"x1.575"x12.657" (WxHxL)
- Vendor: Bel Power

Below are PSU unit drawing supplied by the Bel Power.




Figure 3: PSU Side and Top Views




Figure 4: PSU Front and Back Views




Figure 5: PSU Board Connector View

The table below defines the PSU board connector pins.

Table 3: PSU Output and Control Pin Definition

PIN	NAME	DESCRIPTION
Output		
6, 7, 8, 9, 10	V1	+12 VDC main output
1, 2, 3, 4, 5	PGND	Power ground (return)
Control Pins		
A1	VSB	Standby positive output (+3.3/5 V)
B1	VSB	Standby positive output (+3.3/5 V)
C1	VSB	Standby positive output (+3.3/5 V)
D1	VSB	Standby positive output (+3.3/5 V)
E1	VSB	Standby positive output (+3.3/5 V)
A2	SGND	Signal ground (return)
B2	SGND	Signal ground (return)
C2	HOTSTANDBYEN	Hot standby enable signal

PIN	NAME	DESCRIPTION
D2	VSB_SENSE_R	Standby output negative sense
E2	VSB_SENSE	Standby output positive sense
A3	APS	I ² C address and protocol selection (select by a pull down resistor)
B3	nc	Reserved
C3	SDA	I ² C data signal line
D3	V1_SENSE_R	Main output negative sense
E3	V1_SENSE	Main output positive sense
A4	SCL	I ² C clock signal line
B4	PSON	Power supply on input (connect to A2/B2 to turn unit on)
C4	SMB_ALERT	SMB Alert signal output
D4	nc	Reserved
E4	INOK	DC input OK signal
A5	PSKILL	Power supply kill (lagging pin)
B5	ISHARE	Current share bus (lagging pin)
C5	PWOK	Power OK signal output (lagging pin)
D5	VSB_SEL	Standby voltage selection (lagging pin)
E5	PRESENT_L	Power supply present (lagging pin)

5.2 Fan Modules

Wedge 100BF-64X uses 10 hot swappable fans in the system with 5 for the upper card and 5 for the lower card. 5 fans are mounted on a fan board and the same fan board is used by the upper and the lower card. Wedge 100BF-32X uses 5 hot swappable fans and is equipped with the same fan board.

The fans are 40mm fans with small profile connectors that attaches to the Fan Termination Board without blocking the airflow. Wedge 100B systems support front to back airflow directions. Fans and PSUs have the same airflow direction. Below is a summary of fan features.

- Individually hot swappable fans
- Individually controlled speed with platform software that runs on BMC
- On Power-up, default speed provides enough cooling before software comes up
- Over current protection for each fan module
- Each fan operates individually without interruption from another fan that is failed
- 40 mm fans with 6-15W max power each

Fan connector signals are listed in the following table.

Signal	Description
12V	12V power input to the Fan Module
GND	Ground
FAN_TACH	Tachometer output from fan module for monitoring its speed
FAN_PWM	Pulse Width Modulation signal input to the fan module to set its speed
FAN_PRESENT_L	Fan module present signal. Active Low. Grounded inside the fan module and must be pulled up on the PCB
FAN_ID	Front-to-back or back-to-front airflow direction



6 System Overview

6.1 Main Board

Below are key features on the Wedge 100B main boards.

- Front Panel Ethernet Ports
 - Wedge100BF-65X: 65 QSFP28 ports
 - 16 2x2 zQSFP cage connectors. 8 on the upper board and 8 on the lower board
 - Selective ports on the lower board use repeater device
 - 1 zQSFP connector on the upper board for the 65th port. This port can support all the line rates supported by the other ports with the addition of 1GE.
 - Wedge100BF-32X: 32 QSFP28 ports
 - 8 2x2 zQSFP cage connectors
 - Each front panel port can be configured to support 100GE, 50GE (2x25G), 25GE, 40GE, 4x25GE and 4x10GE (with QSFP to SFP+ breakout cable)
- Barefoot Networks Tofino ASIC
 - Wedge100BF-65X: BFN-T10-064Q device
 - Wedge100BF-32X: BFN-T10-032D device
- BMC (Baseboard Management Controller) for general system control
 - Aspeed AST1250 Microprocessor with 512MB DDR3 memory @ 1.6GHz
 - Tofino power management
 - System voltage monitoring
 - 13 SMBUS
 - Fan control
 - USB 2.0
 - Gigabit Ethernet management port
 - SPI boot flash
 - System reset
 - Interrupts
- Serial console port
- 10/100/1000 Management Port
- USB Type A host port
- System status LEDs
- CPLD for UART muxing, device resets, I2C muxing, device interrupts, LED control
- Temperature sensors
- Power sequence chip for system level voltage sequencing and monitoring
- PEX8717 PCIe Gen3 Switch

6.2 CPU Module

To provide flexibility in system design, a separate CPU module board is used. CPU Module will be COM Express Type 6 which is an industry standard embedded CPU board.

Below is a feature summary of the COM Express board used by Wedge 100BF systems.

- o Intel Xeon D-1517, 4 Core@2.2GHz, 8GB DDR4 SDRAM, 128GB M.2 SSD
- o SPI Bios Flash on the CPU module
- o 12V input to the CPU module
- o 1x GE Ethernet Port (MDI) to GE Switch on Lower Board
- o 2x I2C Interface to BMC on Lower Board
- o 4x USB Interface to Main Switch Board, 1 for Front Panel USB, 2 for CS2112 USB<-> I2C Switches and 1 for the BMC
- o 1x USB Interface to EUSB Flash on Lower Board (I don't think code can be executed directly from this device. The memory module performance needs to be checked)
- o Optional 1xSATA Interface to M.2 SSD Drive or mSata drive on Lower Board
- o 1x UART Interface to CPLD on Lower Board
- o 1x GPO for CPU Module Present to BMC and CPLD on Lower Board
- o 1x GPO for PCIE Reset to CPLD on Lower Board
- o 1x GPI CPU Module Reset
- o PCIe Gen2/3 x4 Interface
- o PCIe Differential Reference Clock to the Lower Board
- o 2x I2C Alert Input to the BMC on Lower Board
- o Temperature Sensor implemented in the CPU Card