



Facebook / Microsoft M.2 Carrier Card Design Specification

Version 1.0

Authors:

Chris Petersen, Hardware System Engineer, Facebook Inc.

Mike Yan, Hardware Engineer, Facebook Inc.

Mark A Shaw, Senior Hardware Engineer Manager, Microsoft Corp

Clark Shao, Hardware Engineer, Facebook Inc.

© 2018 Facebook/Microsoft.



As of Oct 24, 2014, the following persons or entities have made this Specification available under the Open Web Foundation Final Specification Agreement (OWFa 1.0), which is available at:

[http://www.openwebfoundation.org/legal/the-owf-1-0-agreements/owfa-1-0:](http://www.openwebfoundation.org/legal/the-owf-1-0-agreements/owfa-1-0)

Facebook, Inc., Microsoft Corp.

You can review the signed copies of the Open Web Foundation Agreement Version 1.0 for this Specification at

<http://opencompute.org/licensing/>,

which may also include additional parties to those listed above.

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, noninfringement, fitness for a particular purpose, or title, related to the Specification. The entire risk as to implementing or otherwise using the Specification is assumed by the Specification implementer and user. 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.

CONTRIBUTORS AND LICENSORS OF THIS SPECIFICATION MAY HAVE MENTIONED CERTAIN TECHNOLOGIES THAT ARE MERELY REFERENCED WITHIN THIS SPECIFICATION AND NOT LICENSED UNDER THE OWF CLA OR OWFa. THE FOLLOWING IS A LIST OF MERELY REFERENCED TECHNOLOGY: INTELLIGENT PLATFORM MANAGEMENT INTERFACE (IPMI), I2C TRADEMARK OF PHILLIPS SEMICONDUCTOR. IMPLEMENTATION OF THESE TECHNOLOGIES MAY BE SUBJECT TO THEIR OWN LEGAL TERMS.

Table of Contents

1	Overview.....	5
2	Features.....	5
2.1	PCIe.....	5
2.1.1	Bifurcation Settings	5
2.2	Reset	6
2.3	REFCLK	6
2.4	I2C	7
2.5	EEPROM.....	8
2.6	LEDs.....	8
2.7	Power.....	9
2.8	OOB Power and Reset.....	10
2.9	PCB Stackup.....	11
2.10	BOM Options	12
3	BMC	13
3.1	FRUID Contents	13
3.2	Sensors.....	13
3.3	OOB Reset and Power.....	13
4	Connectors.....	14
4.1	PCIe Connector	14
4.2	M.2 Connector.....	17
5	Component Placement and Labeling.....	19
6	Mechanical	20
6.1	Mechanical Drawing.....	20
6.2	M.2 Support	21
6.2.1	M.2 module types	21
6.2.2	M.2 module retention	21

6.2.3	M.2 mounting hole	21
6.3	M.2 Thermal Interface	22
6.3.1	Bottom-side M.2 Thermal Interface.....	22
6.3.2	Top-side M.2 Thermal Interface	22
6.3.3	Thermal Interface Material	22
6.4	Heatsinks.....	24
6.4.1	Top Heat Sink.....	24
6.4.2	Bottom Heat Sink	24
7	Environmental Requirements	25
7.1	Environmental	25
7.2	Thermal Design Guidelines.....	25
7.3	Shock & Vibration	26
7.4	Regulatory.....	26
8	Labels and Markings.....	26
9	Revision History	27

1 Overview

This specification provides the requirements for a PCIe Full Height Half Length (FHHL) form factor card, that supports up to four M.2 form factor solid-state drives (SSDs). The card shall support 110mm (Type 22110) or 80mm (Type 22080) dual sided M.2 modules.

2 Features

2.1 PCIe

The card shall interface to the motherboard through standard PCIe x16 edge card connector. The 16 lanes shall be bifurcated to 4x4 to accommodate the 4 M.2 modules. The card shall be designed to meet the electrical requirements of PCIe Gen3, but can electrically accommodate all generations of PCIe. The lane mapping between the edge connector and the M.2 modules is shown in Figure 1.




Figure 1. PCIe Lane Mapping

2.1.1 Bifurcation Settings

The card shall support 4x4 bifurcation of the x16 PCIe lanes. To support automated bifurcation detection, the bifurcation information is made available to the motherboard and BIOS using pull-up and pulldown resistors on the unused present pin B31. This pin is pulled low on the card to indicate 4x4 bifurcation to the BIOS. Otherwise, bifurcation will need to be set manually in the BIOS or can be read out of band via the BMC from the FRUID EEPROM. The pinout for the PCIe x16 connector is shown in Section 4.1.

2.2 Reset

The card shall support PERST# from the PCIe card edge connector. The signal shall be buffered and fanned out to the individual M.2 modules as shown in Figure 2. The card also supports an optional out of band (OOB) reset enabling reset of individual M.2 modules through an I2C IO expander.

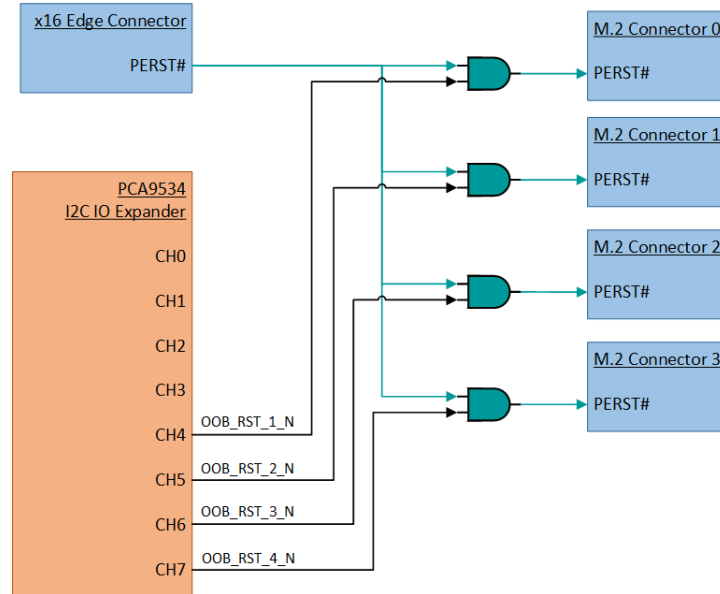


Figure 2. Reset Block Diagram

2.3 REFCLK

The card shall support a single PCIe reference clock (REFCLK) from the PCIe edge card connector. REFCLK shall be buffered and fanned out to the individual M.2 modules using an I2C-configurable clock buffer (IDT 9DB433AGLFT or equivalent) as shown in Figure 3.

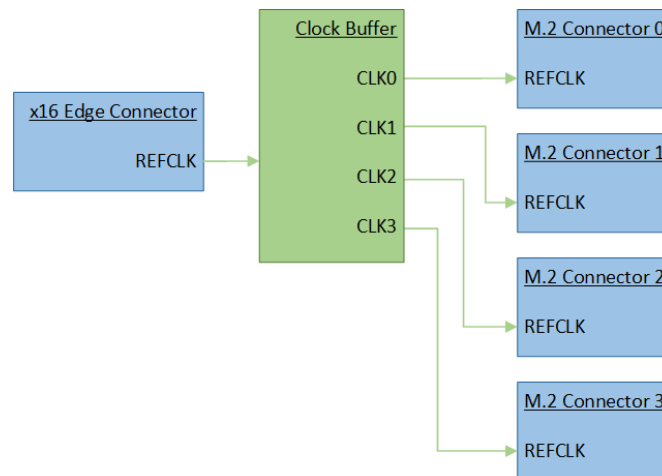


Figure 3. PCIe Clock Block Diagram

2.4 I2C

The card shall support I2C communication with the motherboard BMC through the PCIe card edge connector. The block diagram is shown in Figure 4.

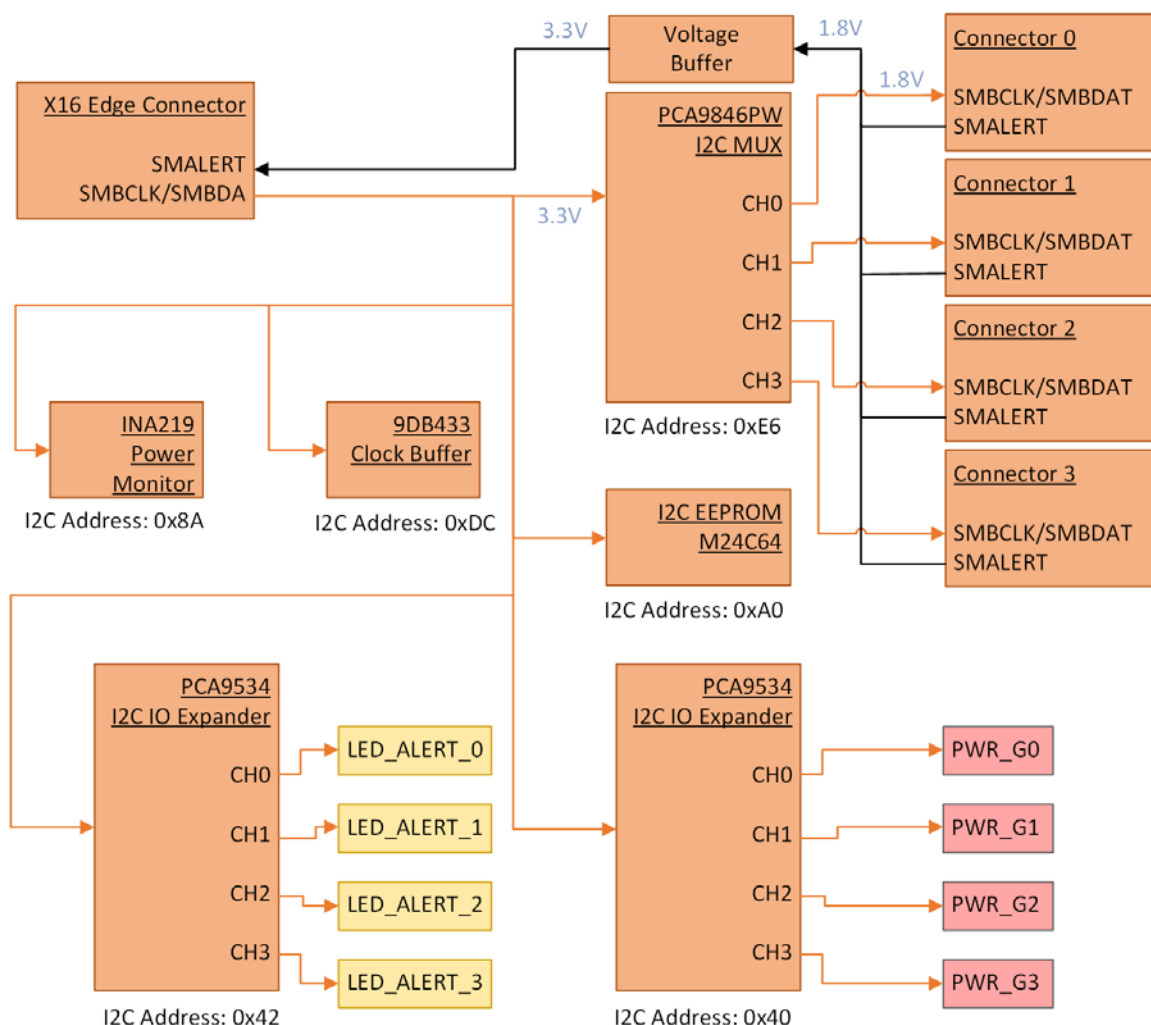


Figure 4. I2C Block Diagram

Pull-ups exist on the host side for SMBCLK, SMBDAT, and SMALERT; however, for the purposes of bring-up, de-populated pull-ups should be provisioned.

The card shall provide the SMBUS connection to the M.2 modules as outlined in the ECN to the PCIe M.2 Specification, titled “SMBus interface for SSD Socket 2 and Socket 3” (August 11, 2014). The M.2 modules shall be connected to the host through an I2C Mux. Note that the logic level required by the M.2 module is 1.8V. Therefore, it is required that the I2C Mux support voltage translation from 3.3V to 1.8V.

For the purposes of bring-up, all I2C address lines should have options for both pull-ups and pull-downs.

2.5 EEPROM

The card shall support a 64Kb EEPROM (FRUID) for storage of manufacturing data. The EEPROM shall be available to the motherboard using the I2C bus.

2.6 LEDs

The card shall support LEDs for communicating state and failure information. Bicolor (red/blue) LEDs should be placed towards the front of the mechanical bracket and should be visible through the front bulkhead. The LED number should be clearly visible on the bracket, along with the silkscreen of its corresponding connector on the PCB. The LED to connector mapping is shown below for both orientations. Silkscreen should be placed on both top and bottom sides of the PCB to clearly indicate LED and connector positions. A block diagram of the LED circuitry is shown in Figure 5.

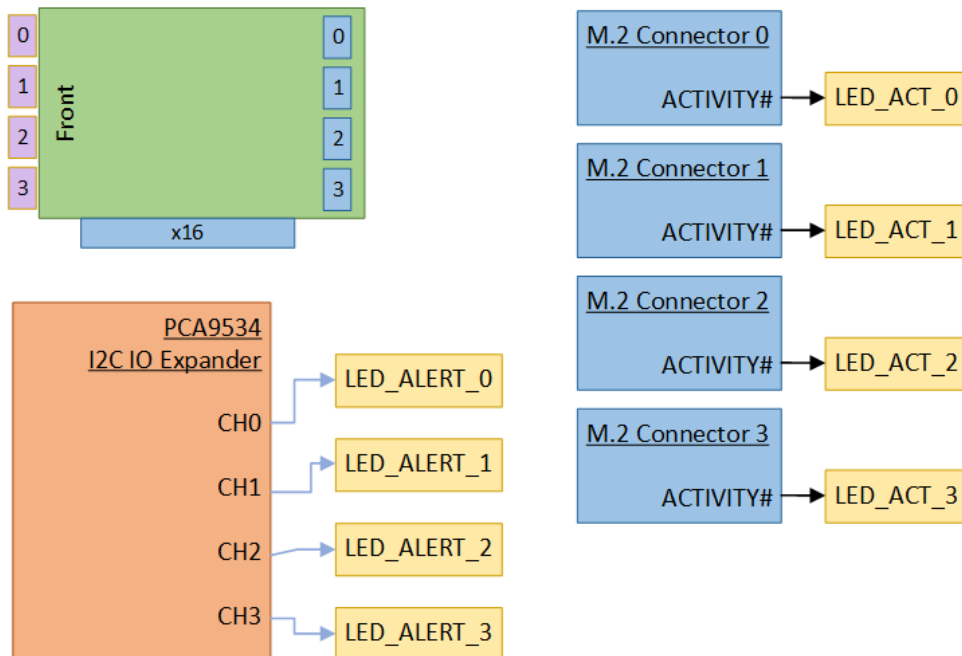


Figure 5. LED Block Diagram

The blue LEDs are to indicate M.2 activity, and should be driven by the LED signal (Pin 10) of the M.2 modules. The LED should be on when the SSD is installed and idle, and flashing when there is activity.

The red LEDs are to indicate module fault, and should be driven by the I2C-addressable GPIO expander.

2.7 Power

The card shall support 12V, 3.3V, and 3.3V Aux power through the PCIe edge connector. A block diagram of the power tree is shown in Figure 6.

12V shall be supplied to the card through a power monitor that allows for voltage, current, and power reporting over the I2C bus. As an option, a hot-swap controller can be provided on this rail, which limits the in-rush current and isolates the rest of the system from any power fault on the card. Power to the M.2 modules shall be generated from a 3.3V switching converter that is supplied by the 12V output of the hot swap controller. The converter shall be scaled to support 8.5W per M.2 module up to 14.85W per M.2 module.

Optionally, the card shall support FETs to enable the system to control power to individual M.2 modules. This enables the system to clear errors by cycling power to specific modules. The card shall contain the necessary design components to bypass this feature if it is not required by the system.

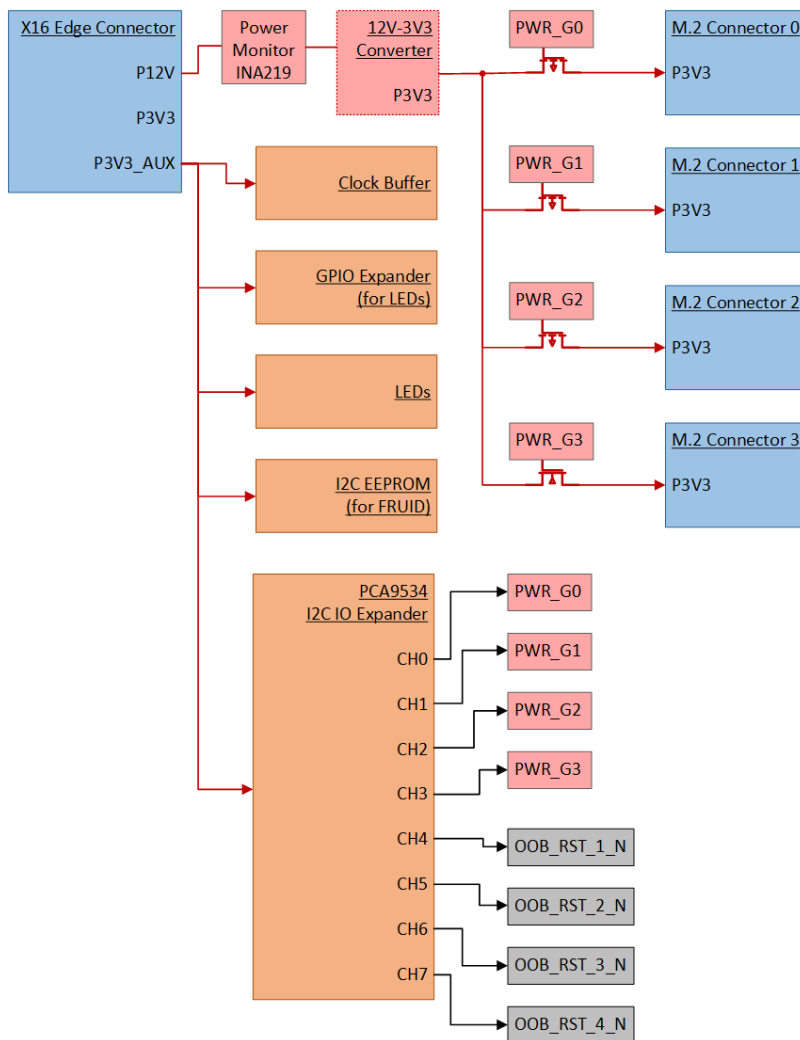


Figure 6. Power Block Diagram

2.8 OOB Power and Reset

Timing for control of the 3.3V power enable and reset of the M.2 modules should adhere to the PCIe specification. Figure 7 shows the recommended timing.

- PERST_RISER_N - Active low signal coming from the motherboard (PERST#).
- OOB_RST_1_N is the active low signal coming from the I2C IO expander shown in Figure 2.
- PERST_M2_1_N is the PERST# input to the M.2 module.
- PWR_G1 is the FET control signal driven by the I2C IO expander shown in Figure 6
- P3V3_M2_1 is the 3.3V power output of the FET to the M.2 module




Figure 7. PCIe reset timing diagram for optional FET control

The timing sequence is as follows:

- T0-T1: Normal operation.
- T1: To initiate a reset, the GPIO OOB_RST_1_N is asserted.
- T1-T2: After sufficient timing as per the PCIe specification, PWR_G1 is de-asserted which causes P3V3_M2_1 to fall out of regulation.
- T3: After the module has been powered off for a sufficient amount of time, PWR_G1 is asserted.
- T4: P3V3_M2_1 is stable once again.
- T4-T5: After sufficient timing as per the PCIe specification, OOB_RST_1_N is released.

2.9 PCB Stackup

The stackup should use a sufficiently low-loss material, such that the PCIe signals can be driven at 8 GT/s without the need for re-drivers. The carrier card plus the M.2s must meet the channel loss requirements as defined in the PCIe CEM specification.

The PCB should have an ENIG finish or equivalent to prevent oxidation.

To enable routing of high speed traces on inner layers, an 8-layer stackup is suggested. An example is shown below. It is possible to reduce the layer count through layout optimizations.

Table 1. The M.2 Carrier Card 8 Layers Stackup

STACKUP		Target Z (ohms) - MicroStrip				50	Breakout	85		100		Breakout	
		Target Z (ohms) - StripLine				50		85		100			
		Z tolerance				±10%		±10%		±10%			
		Z Type				Single		Differential		Differential		Differential	
Layer#	Description	Copper Weight (oz)	Thickness	Tolerance	Er	Width	Width	Width	Space	Width	Space	Width	Space
	Soldermask		0.60		3.8								
1	TOP	0.5+plating	1.95			6.1	4	5.3	4	5	8	3.9	4
	PP		3.70	±0.709	3.9								
2	GND	1	1.30										
	CORE		4.00	±0.709	3.9								
3	IN1	1	1.30			4.7	4	5	6.1	4.3	12	3.9	4
	PP		14.00	±1.97	4								
4	VCC	2	2.60										
	CORE		4.10	±0.709	3.9								
5	VCC1	2	2.60										
	PP		14.00	±1.97	4								
6	IN2	1	1.30			4.7	4	5	6.1	4.3	12	3.9	4
	CORE		4.00	±0.709	3.9								
7	GND1	1	1.30										
	PP		3.70	±0.709	3.9								
8	BOTTOM	0.5+plating	1.95			6.1	4	5.3	4	5	8	3.9	4
	Soldermask		0.60		3.8								
		Total	63.00	±10%									

2.10 BOM Options

The following table describes 2x BOM options, but additional combinations are possible:

Table 2. The M.2 Carrier Card BOM Options

#	Description	SKU 1	SKU 2	Notes
1	Use Hot-swap Controller	Yes	No	For power protection
2	SMBUS lines connected to 3.3V instead of 1.8V	No	No	For future proofing
3	High power 12V-3.3V VR (18A @ 3.3V)	Yes	No	For future proofing
4	Low power VR (10.5A @ 3.3V)	No	Yes	
5	Bypass M.2 power-control FETs	No	No	
6	Populate attention LEDs	Yes	Yes	
7	Populate status LEDs	Yes	Yes	

3 BMC

3.1 FRUID Contents

The format of the contents of the FRUID EEPROM should follow the IPMI Platform Management FRU Information Storage Definition v1.0.

The contents of the FRUID are defined in the table below. The table only defines the minimum number of required bytes, but if longer fields are needed adding additional bytes is acceptable.

Table 3. FRUID Contents

Name	Minimum byte count	Value
Board Mfg Date	3	Date the board was manufactured
Board Mfg	6	Manufacturer
Board Product	19	“Ava-M.2-SSD-Adapter”
Board Serial Number	11	Supplier defined serial number
Board Part Number	11	Supplier defined part number
Board FRU ID Version	2	Start at 0.1 and increment with every change
Board version	3	One of the following: “EVT”, “DVT”, “PVT”, “MP”
LR_ID_STRING	34	“x4x4x4x4 PCIe 3.0 M.2 carrier card”

3.2 Sensors

The BMC should support reading of the thermal sensors to drive the FSC, as well as the voltage/current/power readings from the input P12V monitor.

3.3 OOB Reset and Power

The BMC should support the ability to perform OOB reset and power functionality as described in Section 2.8.

4 Connectors

4.1 PCIe Connector

The blade shall support a standard PCIe x16 edge card connector. The pinout for supporting PCIe x16 described in Table 4. For further information, refer to the PCI Express® Card Electromechanical Specification. Note that pin B31 has been repurposed to support automated bifurcation. Additionally, pin B12 has been repurposed to add a SMBus Alert signal.

The card shall also implement the hot-plug present detection scheme as defined in the PCIe CEM spec. Since this is a x16 card, pin B81 (PRSNT#2) should be connected to pin A1 (PRSNT#1).

Table 4. PCIe x16 Connector Pinout

Pin	Side B Connector		Side A Connector	
#	Name	Description	Name	Description
1	+12v	+12 volt power	PRSNT#1	Hot plug presence detect
2	+12v	+12 volt power	+12v	+12 volt power
3	+12v	+12 volt power	+12v	+12 volt power
4	GND	Ground	GND	Ground
5	SMCLK	SMBus clock	JTAG2	TCK
6	SMDAT	SMBus data	JTAG3	TDI
7	GND	Ground	JTAG4	TDO
8	+3.3v	+3.3 volt power	JTAG5	TMS
9	JTAG1	+TRST#	+3.3v	+3.3 volt power
10	3.3Vaux	3.3v volt power	+3.3v	+3.3 volt power
11	WAKE#	Link Reactivation	PERST#	Fundamental reset
Mechanical Key				
12	SMBALERT#	SMBus Alert	GND	Ground
13	GND	Ground	REFCLK+	Reference Clock
14	PETP(0)	Transmitter Lane 0,	REFCLK-	Differential pair
15	PETN(0)	Differential pair	GND	Ground
16	GND	Ground	PERP(0)	Receiver Lane 0,
17	PRSNT#2	Presence detect	PERN(0)	Differential pair
18	GND	Ground	GND	Ground
19	PETP(1)	Transmitter Lane 1,	RSVD	Reserved
20	PETN(1)	Differential pair	GND	Ground
21	GND	Ground	PERP(1)	Receiver Lane 1,
22	GND	Ground	PERN(1)	Differential pair

23	PETP(2)	Transmitter Lane 2,	GND	Ground
24	PETN(2)	Differential pair	GND	Ground
25	GND	Ground	PERP(2)	Receiver Lane 2,
26	GND	Ground	PERN(2)	Differential pair
27	PETP(3)	Transmitter Lane 3,	GND	Ground
28	PETN(3)	Differential pair	GND	Ground
29	GND	Ground	PERP(3)	Receiver Lane 3,
30	RSVD	Reserved	PERN(3)	Differential pair
31	BIFURx4	0 = PCIe x4 Bifurcation	GND	Ground
32	GND	Ground	RSVD	Reserved
33	PETP(4)	Transmitter Lane 4,	RSVD	Reserved
34	PETN(4)	Differential pair	GND	Ground
35	GND	Ground	PERP(4)	Receiver Lane 4,
36	GND	Ground	PERN(4)	Differential pair
37	PETP(5)	Transmitter Lane 5,	GND	Ground
38	PETN(5)	Differential pair	GND	Ground
39	GND	Ground	PERP(5)	Receiver Lane 5,
40	GND	Ground	PERN(5)	Differential pair
41	PETP(6)	Transmitter Lane 6,	GND	Ground
42	PETN(6)	Differential pair	GND	Ground
43	GND	Ground	PERP(6)	Receiver Lane 6,
44	GND	Ground	PERN(6)	Differential pair
45	PETP(7)	Transmitter Lane 7,	GND	Ground
46	PETN(7)	Differential pair	GND	Ground
47	GND	Ground	PERP(7)	Receiver Lane 7,
48	PRSNT#2	Hot plug detect	PERN(7)	Differential pair
49	GND	Ground	GND	Ground
50	PETP(8)	Transmitter Lane 8,	RSVD	Reserved
51	PETN(8)	Differential pair	GND	Ground
52	GND	Ground	PERP(8)	Receiver Lane 8,
53	GND	Ground	PERN(8)	Differential pair
54	PETP(9)	Transmitter Lane 9,	GND	Ground
55	PETN(9)	Differential pair	GND	Ground
56	GND	Ground	PERP(9)	Receiver Lane 9,
57	GND	Ground	PERN(9)	Differential pair
58	PETP(10)	Transmitter Lane 10,	GND	Ground

59	PETN(10)	Differential pair	GND	Ground
60	GND	Ground	PERP(10)	Receiver Lane 10, Differential pair
61	GND	Ground	PERN(10)	
62	PETP(11)	Transmitter Lane 11,	GND	Ground
63	PETN(11)	Differential pair	GND	Ground
64	GND	Ground	PERP(11)	Receiver Lane 11, Differential pair
65	GND	Ground	PERN(11)	
66	PETP(12)	Transmitter Lane 12,	GND	Ground
67	PETN(12)	Differential pair	GND	Ground
68	GND	Ground	PERP(12)	Receiver Lane 12, Differential pair
69	GND	Ground	PERN(12)	
70	PETP(13)	Transmitter Lane 13,	GND	Ground
71	PETN(13)	Differential pair	GND	Ground
72	GND	Ground	PERP(13)	Receiver Lane 13, Differential pair
73	GND	Ground	PERN(13)	
74	PETP(14)	Transmitter Lane 14,	GND	Ground
75	PETN(14)	Differential pair	GND	Ground
76	GND	Ground	PERP(14)	Receiver Lane 14, Differential pair
77	GND	Ground	PERN(14)	
78	PETP(15)	Transmitter Lane 15,	GND	Ground
79	PETN(15)	Differential pair	GND	Ground
80	GND	Ground	PERP(15)	Receiver Lane 15, Differential pair
81	PRSNT#2	Hot plug present detect	PERN(15)	
82	RSVD#2	Hot Plug Detect	GND	Ground

4.2 M.2 Connector

The recommended M.2 connector shall have a height of 5.8mm to enable the use of thermal pads above and below the M.2 to aid in cooling. A list of available part numbers is shown in table 5.

Table 5. M.2 Connector Part Numbers

Supplier	Connector Part Number
Amphenol	MDT580M02001
Arogsy	NASM0-S6730-TS58

The pin-out shall follow the M.2 Socket 3, M-key pin-out defined in the M.2 specification. The pinout includes SMBus pin definitions contained in the ECN to the PCIe M.2 Specification, titled “SMBus interface for SSD Socket 2 and Socket 3” (August 11, 2014). This update assigns SMBus pins to pins 44, 42, and 40. The pinout is shown in table 6.

Table 6. M.2 Connector Pinout

M.2 Module Standard Pinout					
Pin	Signal	Description	Pin	Signal	Description
74	3.3V	3.3V Power.	75	GND	Ground
72	3.3V	3.3V Power	73	GND	Ground
70	3.3V	3.3V Power	71	GND	Ground
68	SUSCLK(32KHz)	Reduce Power Clock	69	NC	Reserved
66	KEY	Module Key	67	NC	Reserved
64	KEY	Module Key	65	KEY	Module Key
62	KEY	Module Key	63	KEY	Module Key
60	KEY	Module Key	61	KEY	Module Key
58	NC	Reserved	59	KEY	Module Key
56	NC	Reserved	57	GND	Ground
54	PEWAKE#	PCIe PME Wake (OD)	55	REFLKCP	PCIe Reference Clock
52	CLKREQ#	Reference Clock Request	53	REFLKN	PCIe Reference Clock
50	PERST#	PCIe Reset	51	GND	Ground
48	NC	Reserved	49	PETP0	PCIe Transmit Lane 0
46	NC	Reserved	47	PETN0	PCIe Transmit Lane 0
44	SMBALERT#	SMBus Alert	45	GND	Ground
42	SMDATA	SMBus Data	43	PERP0	PCIe Receive Lane 0
40	SMCLK	SMBus Clock	41	PERN0	PCIe Receive Lane 0
38	DEVSLP	Device Sleep	39	GND	Ground
36	NC	Reserved	37	PETP1	PCIe Transmit Lane 1

34	NC	Reserved	35	PETN1	PCIe Transmit Lane 1
32	NC	Reserved	33	GND	Ground
30	NC	Reserved	31	PERP1	PCIe Receive Lane 1
28	NC	Reserved	29	PERN1	PCIe Receive Lane 1
26	NC	Reserved	27	GND	Ground
24	NC	Reserved	25	PETP2	PCIe Transmit Lane 2
22	NC	Reserved	23	PETN2	PCIe Transmit Lane 2
20	NC	Reserved	21	GND	Ground
18	3.3V	3.3V Power	19	PERP2	PCIe Receive Lane 2
16	3.3V	3.3V Power	17	PERN2	PCIe Receive Lane 2
14	3.3V	3.3V Power	15	GND	Ground
12	3.3V	3.3V Power	13	PETP3	PCIe Transmit Lane 3
10	DAS/DSS#	Drive Active Signal (OD)	11	PETN3	PCIe Transmit Lane 3
8	NC	Reserved	9	GND	Ground
6	NC	Reserved	7	PERP3	PCIe Receive Lane 3
4	3.3V	3.3V Power	5	PERN3	PCIe Receive Lane 3
2	3.3V	3.3V Power	3	GND	Ground
			1	GND	Ground

5 Component Placement and Labeling

The M.2 connectors shall be silkscreened in large font, on both top and bottom sides of the card.

No silkscreen shall be placed over high speed traces on the top or bottom layer.

6 Mechanical

6.1 Mechanical Drawing

The card shall meet the PCIe CEM Specification Rev 3.0 for a full height, half-length form factor PCIe card including top and bottom side component height and keepout requirements. M.2 modules shall be oriented to support a longitudinal airflow allowing air to traverse the length of the card before intersecting with the M.2 connectors. The mechanical drawing is shown in Figure 8.




Figure 8. Mechanical Drawing

6.2 M.2 Support

6.2.1 M.2 module types

The card shall support type 2280 and 22110 M.2 modules. It must also support M.2 module with a maximum of 1.5mm bottom-side height and a maximum top-side height of 3.1mm.

6.2.2 M.2 module retention

The M.2 module retention is a single piece that presses into the board, and retains the modules with a plastic tab. An example shown in Figure 9.




Figure 9. M.2 module retention

6.2.3 M.2 mounting hole

The M.2 retention clip requires a hole in the PCB with the dimensions as shown in Figure 10.




Figure 10 M.2 mounting hole dimensions

6.3 M.2 Thermal Interface

The card has various features to facilitate the conduction of heat away from both the top- and bottom-side of the M.2 modules.

6.3.1 Bottom-side M.2 Thermal Interface

The card shall support a keepout area under the M.2 modules as outlined in Figure 11 and Figure 12 on next page. The keepout is defined such that a solid copper pour with the same shape is present on the top layer. The intention is to use this copper pour to help conduct heat from the M.2 modules into the PCB through a thermal interface material (TIM).

To increase the efficacy of this thermal solution, copper pours of the same shape shall exist on all layers, to the extent that routing permits. These copper pours shall all be tied together with large thermal stitching vias to the solid GND copper planes on the internal layers of the board. This will help increase the through-plane conductivity of the board, to assist heat transfer from the top-side of the card to the bottom-side.

Furthermore, to increase the heat dissipation from the bottom-side of the card, a heatsink shall be installed on the bottom-side with the mounting holes described in the mechanical drawing. A similar TIM shall be used to assist heat transfer from the bottom-side of the card to the heatsink. This bottom-side heatsink assembly shall be designed such that the overall profile still adheres to the PCIe CEM dimensions. It follows that this bottom-side area shall be free of all components.

Because the TIM will have a different dielectric constant from air, high speed traces are not to be routed such that they are in a layer that touches the TIM.

6.3.2 Top-side M.2 Thermal Interface

Similarly, a combination of TIM and heatsink can be implemented on the top-side of the card to improve thermal performance. Here, a TIM is installed between the top-side of the M.2 module and a heatsink, which clips to the top-side of the card with a tool-less latch design.

6.3.3 Thermal Interface Material

In order to accommodate different M.2 module layouts with varying component heights, thermal pads (also known as gap pads) are recommended. Selected gap pads should meet the following requirements:

- Thermal pads should be inherently soft and sized appropriately, to ensure bending in the carrier card and M.2 module PCB meet appropriate requirements or specifications
- Bulk thermal conductivity of selected gap pads should maintain M.2 modules within recommended temperatures under all operating conditions
- Gap pads should not leave significant residue on the carrier card and top heat sinks to minimize time required to service or replace M.2 modules




Figure 11. Component Keepout - Top Side




Figure 12. Component Keepout – Bottom Side

6.4 Heatsinks

The carrier card shall support the following heatsinks to permit dissipation of heat from both sides of the card.

6.4.1 Top Heat Sink

The top heat sink is responsible for dissipating heat from the top side of each M.2 module. Fin count and density should be selected such that cooling requirements are met with minimal impact on airflow impedance. It should be a tool-less design and separate for each M.2 module to accommodate the use of soft gap pads with minimal rebound. It should also permit each M.2 module on the carrier card, and its corresponding retention clip, to be accessed separately and with ease. When installed, the top heat sink should meet PCIe CEM dimensional requirements.

6.4.2 Bottom Heat Sink

The bottom heat sink is responsible for dissipating heat from the carrier card and incorporates thermal pads to accommodate the same. It should include mounting features, corresponding to the mounting holes in the carrier card, which ensure the thermal pads make adequate contact with the PCB. When installed, the bottom heat sink should meet PCIe CEM dimensional requirements.

7 Environmental Requirements

7.1 Environmental

This carrier card shall meet the same environmental requirements specified at the system level as below.

- Gaseous contamination: Severity Level G1 per ANSI/ISA 71.04-1985
- Ambient operating temperature range: -5°C to +45°C
- Storage temperature range: -40°C to +70°C (long-term storage)
- Transportation temperature range: -55°C to +85°C (short-term storage)
- Operating and storage relative humidity: 10% to 90% (non-condensing)
- Operating altitude up to 6,000 feet

7.2 Thermal Design Guidelines

The information in this section is intended to serve as a quick reference for the M.2 carrier card adopters. It should be used as reference for design and feasibility, and should not replace detailed thermal analysis. The system adopting the M.2 carrier card should drive the appropriate TIM material and heatsink design/selection to maintain appropriate thermal margins for components on the M.2's.

This guideline is based on M.2 modules, with a maximum power of 7.9W each, in a wind tunnel enclosure based on PCI Express standards for a full height, half-length device. Airflow is specified across the card (direction from IO bracket to M.2 connectors) and the approach cross section is 114.5mm by 20.04mm. The curve showed in Figure 13 was obtained by CFD analysis of a reference design. Additional geometric parameters and material properties used in this simulation are summarized in table 7.




Figure 13. Airflow requirement for the M.2 carrier card cooling

Table 7. Reference for the M.2 carrier card geometry and boundaries

The M.2 Carrier Card Reference Design	
Heatsink Type	Pin Fin
Heatsink Width [mm]	16.2
Heatsink Length [mm]	96
Heatsink Height [mm]	4
Heatsink Base Thickness [mm]	1.5
Heatsink Conductivity [W/m-K]	96
TIM Conductivity [W/m-K]	2.7

7.3 Shock & Vibration

This carrier card shall meet the same shock and vibration requirements specified at the system level as shown in table 8.

Table 8. Shock and Vibration Requirements

	Operating	Non-Operating
Vibration	0.5g acceleration, 5 to 500 Hz, 10 sweeps at 1 octave/minute per each of the three axes (one sweep is 5 to 500 to 5 Hz)	1g acceleration, 5 to 500 Hz, 10 sweeps at 1 octave/minute per each of the three axes (one sweep is 5 to 500 to 5 Hz)
Shock	6g, half-sine 11mS, 5 shocks per each of the three axes	12g, half-sine 11mS, 10 shocks per each of the three axes

7.4 Regulatory

This carrier card shall meet CE, CB, FCC Class A, WEEE, and ROHS requirements.

8 Labels and Markings

The carrier card shall have the following barcoded labels in visible locations where they can be easily scanned during integration. Vendor and Facebook will have an agreement for the label locations.

- Vendor P/N, S/N, REV (revision would increment for any approved changes)
- Facebook P/N (or OCP customer P/N)
- Date code (industry standard: WEEK/YEAR)
- The assembly shall be marked “THIS SIDE UP”, “TOP SIDE”, “UP ^” or other approved marking in bright, large characters in a color to be defined by ODM and Facebook (or OCP customer). This printing may be on the PCB itself, or on an installed component such as an air baffle. The label should be clear and easy to read in low light conditions, when viewed from above or below from 2 feet away and at an angle of approximately 60 degrees off horizontal.

9 Revision History

Author	Description	Revision	Date
Chris Petersen	Initial draft spec	0.1	1/29/2016
Mark / Chris	Lots of changes	0.2	3/1/2016
Mark / Chris	1 st OCP release	0.3	4/11/2016
Chris	<ul style="list-style-type: none"> - Added FRU information - M.2 connector information - Fixed power block diagram to place clock buffer on P3V3_AUX - Fixed some typos in the pinouts 	0.4	5/25/2016
Chris	<ul style="list-style-type: none"> - Moved hot-swap controller as optional, while keeping power monitoring as a requirement - Added 8-layer stackup guidance - Added description of thermal solution - Clarified description of bifurcation setting pull-up / pull-down (table remains unchanged) 	0.5	6/3/2016
Mike Yan	Clean-up for OCP official release <ul style="list-style-type: none"> - Updated mechanical and thermal portion to reflect latest design - Updated environmental requirements 	0.6	6/29/2017
Clark Shao	Updated some tables' name and change revision from v0.6 to v1.0 Added section 7.2 for thermal design reference	1.0	3/28/2018