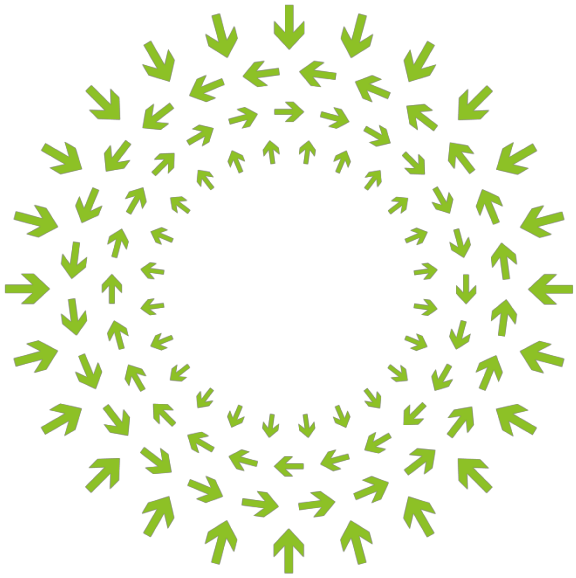

OPEN

Compute Project

M.2 Accelerator Module Hardware Specification

V0.1

Authors:

Hao Shen, Hardware, Facebook

Chris Petersen, Hardware, Facebook

Pavan Shetty, Hardware**,** Facebook

Yueming Li, Thermal, Facebook

Ashwin Narasimha, RTP, Facebook

Harsha Bojja, MQE, Facebook

Ben Wei, Software, Facebook

Sam Naghshineh, TPM, Facebook

1 Scope

M.2 is a standard format defined by PCI-SIG. This specification defines an implementation of PCIe-attached accelerators in a M.2 form factor for use in datacenter environments.

The purpose of this document is to the define module level hardware specification.

2 Contents

3 Overview

The high-level hardware specifications are listed below:

Table 1. Hardware Specification Table

Outline	22 x 110mm
Pin Definition	M.2 22110 Socket 3 key M Pin Definition with some NC pins redefined.
Board Thickness	0.8mm +- 10%
Board Layer Count	Maximum 12 layers
Component Height	2.5mm on top layer and 1.5mm on bottom layer[1]
Power	10W sustained RMS power, refer to chapter 5.2 for more details.
PCIE	Gen4 or Gen3, 4 lanes, refer to chapter 5.4 for more details

1. Height keepout just defines SMT component. It does not include the heat sink or thermal materials.
2. PCIE shall be compliance to PCI Express Base Specification Revision 3.1/4.0

The module includes one ASIC, 2x 64bit LPDDR4x DRAMs, and supporting circuits. All of the circuitry shall be within the M.2 standard form factor and should not require any auxiliary circuit on the system mother board to meet the interface spec.

The recommended placement of M.2 module is shown in Fig.1. The ASIC shall be placed on the top side of the board and ideally is expected to be sitting in the center area of the board for optimal thermal performance. The supplier can adjust the placement based on the tradeoff between thermal and routing study.

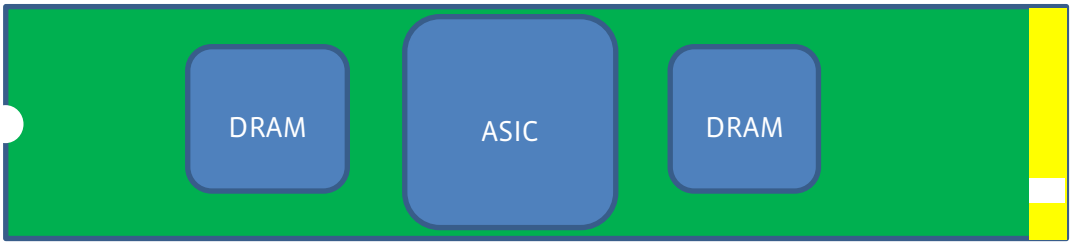

Figure 1. M.2 key component placement drawing

A module is expected to include, but not limited to the following components:

1. ASIC to run the workload
2. LPDDR4x Chips
3. Power Circuits
4. SMBus EEPROM or on-chip memory to save FRU information, 8bit SMBus address 0xA6 as defined in NVMe Specification.

5. Storage for boot firmware if needed but NAND Flash is not allowed.

4 System Block Diagram

The M.2 Accelerator modules will be plugged in to the Glacier Point V2 (GPv2) card, which is a PCIE extension card defined in the Yosemite V2 system.

Glacier Point V2 card is the carrier card inserted in the Yosemite V2 sled. It is plugged in slot 1 and 3 in the below graph. It pairs with the twin lake server to form a subsystem. Slot1 and slot2 are a pair and slot 3 and slot 4 are another pair.

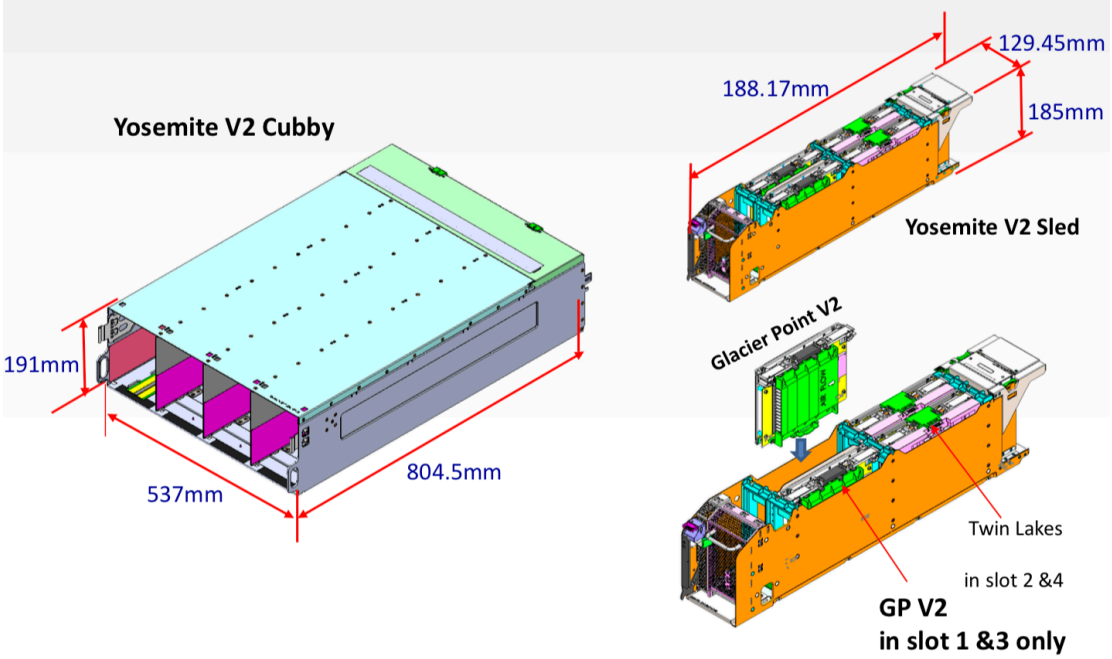

Figure 2. Yosemite V2 and Glacier Point V2 System Drawing

Previously we have designed a Glacier Point V1 card which is a M.2 carrier card in YV2 system which supports 6x M.2 cards. In the GPv2 design, we have made several changes to better support accelerator modules. The major changes are:

- Each GPv2 card can support up to 12x M.2 cards
- Added a PCIe Switch to fanout PCIe lanes
- Included a bridge micro-controller to manage the sideband of each M.2 card
- Included a CPLD to mux UART and JTAG of each module to the debug interface
- Added a power switch to each M.2 card to allow control software to perform complete power cycle of each M.2 cards independently when system is in powered on

PCIe block diagram is listed in Figure 3. There is a Microsemi PM8535 PCIe switch to fanout 24 PCIe lanes from USP to support 12 PCIe Gen3 x4 links at DSP.

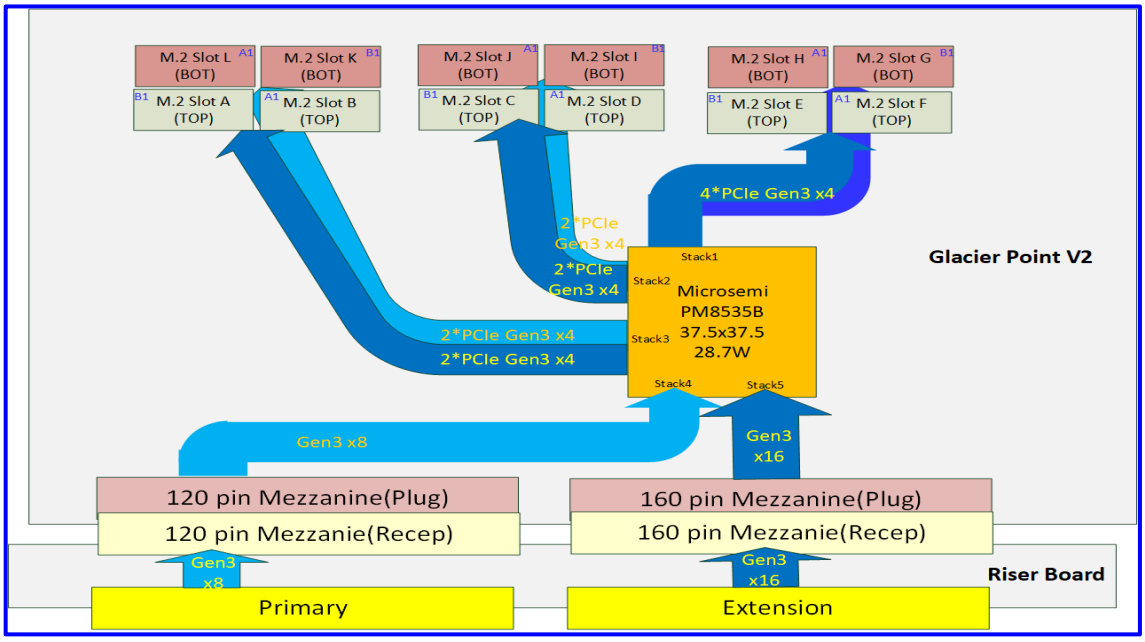

Figure 3. PCIe Block Diagram on GPv2

The block diagram of the SMBus topology in the GPv2 system is listed below. M.2_A through L represent the M.2 connectors. In this graph, SNOWFLAKE represents the bridge IC, which is a micro-controller that manages the GPv2 board. INA231 is the voltage and current monitor. PCA9846PW is the SMBus Mux.

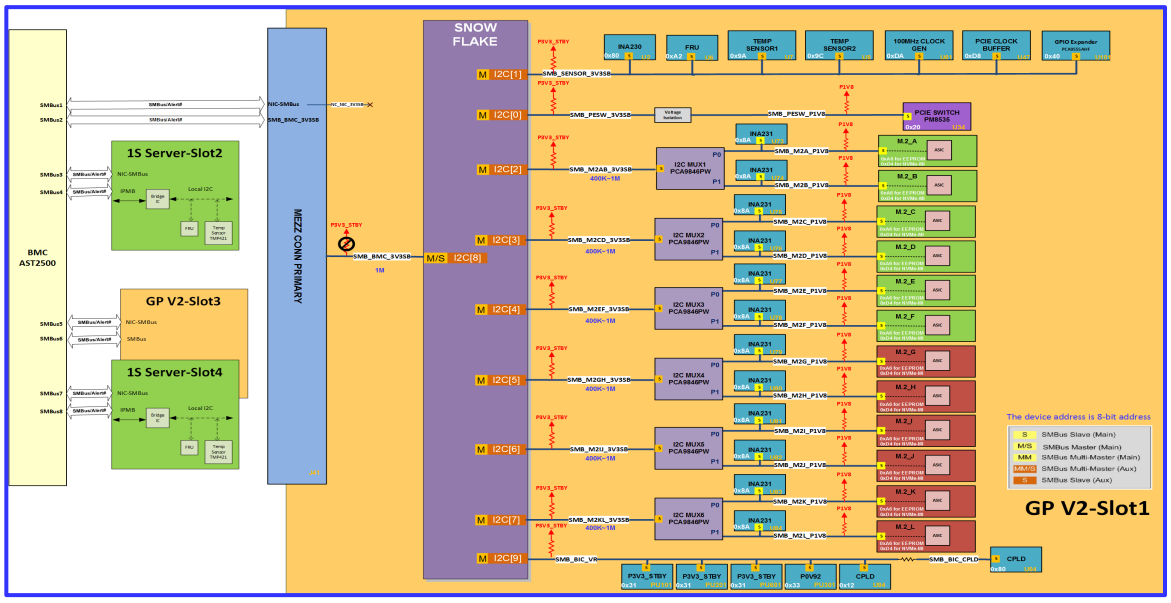

Figure 4. Typical GPv2 SMBus/I2C topology

Yosemite V2 design has defined the debug port on the front panel to ensure the accessibility. Meanwhile the remote debug capability is a feature that is useful in the fleet so that we can dump the error log once the system fails without the need to have the operator to be personally there.

In GPv2 we design with common clock topology. The host CPU will provide the clock to all the M.2 modules through a clock buffer on GPv2 card. PM8535 PCIE switch provide PCIE clock to all M.2 connectors.

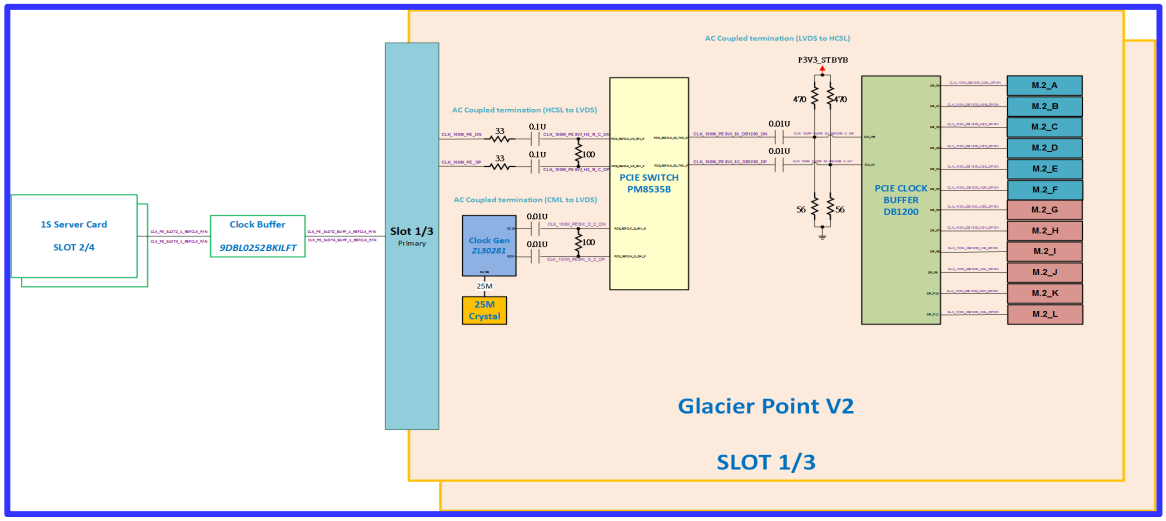

Figure 5: PCIE clock tree

JTAG and UART interfaces are muxed to the system through a CPLD chip. UART port will finally connect both BMC chip and the debug port on the front of the YV2 baseboard. The block diagram is listed below:

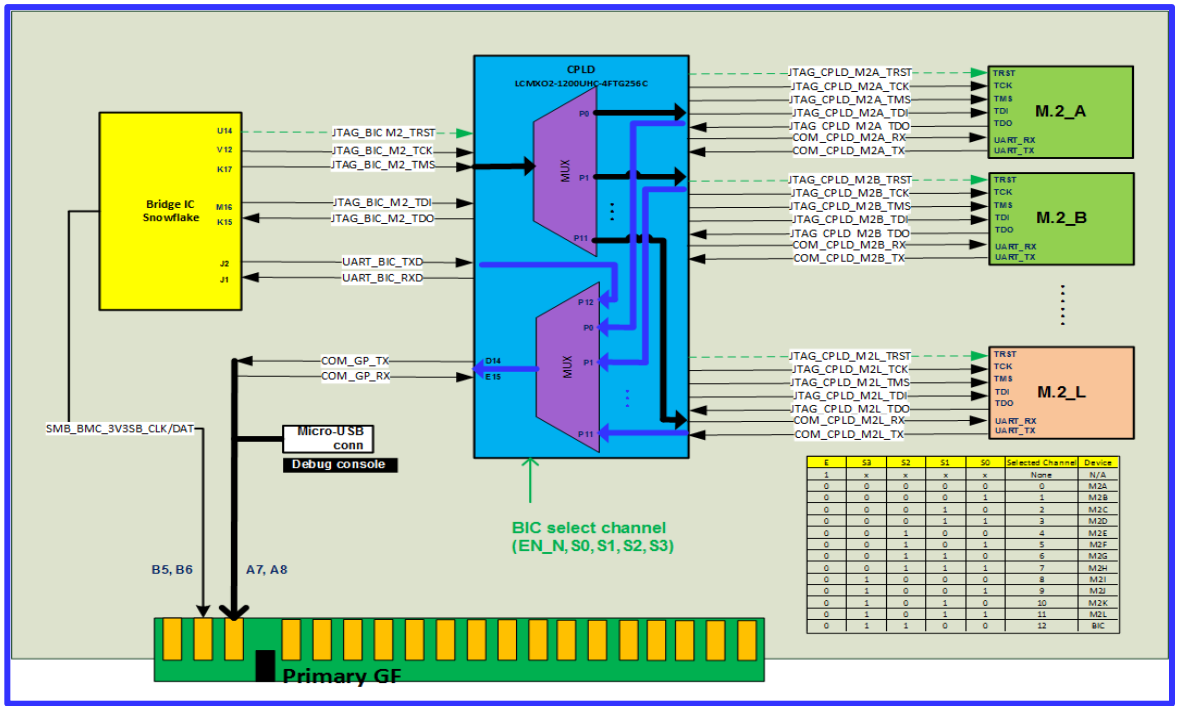

Figure 6. JTAG and UART block on GPv2 board

5 Module Hardware Spec

5.1 Pin Definition

M.2			
Signal	Pin	Pin	Signal
3.3V	2	1	GND
3.3V	4	3	GND
PWRDIS	6	5	PETn3
PLN#	8	7	PETp3
LED_1#(0)	10	9	GND
3.3V	12	11	PERn3
3.3V	14	13	PERp3
3.3V	16	15	GND
3.3V	18	17	PETn2
TRST	20	19	PETp2
VIO_1V8	22	21	GND
TDI	24	23	PERn2
TDO	26	25	PERp2
TCK	28	27	GND
PLA_S3#	30	29	PETn1
GND	32	31	PETp1
USB_D+	34	33	GND
USB_D-	36	35	PERn1
GND	38	37	PERp1
SMB CLK (I/O)(0/1.8V)	40	39	GND
SMB DATA (I/O)(0/1.8V)	42	41	PETn0
ALERT# (O)(0/1.8V)	44	43	PETp0
Reserved_UART_Rx	46	45	GND
Reserved_UART_Tx	48	47	PERn0
PERST# (I)(0/3.3V)	50	49	PERp0
CLKREQ# (I/O)(0/3.3V)	52	51	GND
PEWAKE#(I/O)(0/3.3V)	54	53	REFCLKn
Reserved for MFG DATA	56	55	REFCLKp
Reserved for MFG CLOCK	58	57	GND
ADD IN CARD KEY M			ADD IN CARD KEY M
ADD IN CARD KEY M			ADD IN CARD KEY M
ADD IN CARD KEY M			ADD IN CARD KEY M
ADD IN CARD KEY M			ADD IN CARD KEY M
NC	68	67	TMS
3.3V	70	69	NC
3.3V	72	71	GND
3.3V	74	73	VIO_CFG_GND
		75	GND

Figure 7. M.2 Module Pinout Table

This pinout table and I/O direction is defined from the perspective of the module (not the baseboard perspective). Pin definition is compatible to PCI-SIG M.2 specification though we have re-defined several NC pins mainly for debug purposes. These features are expected to improve the debug capability once hardware is deployed in large scale data center environment.

Table 2. M.2 Module Pinout description

Interface	Signal Name	I/O	Description	Voltage	Requirement
Power Ground	3.3V(9 pins)	I	3.3V running power source	3.3V	Required
	GND(15 pins)		Ground	0V	Required
	VIO_1V8	I	Reserved 1.8V running power source for future PCI-SIG standard.	1.8V	NC in module, the platform will leave this pin open.
PCIe	PETp0/PETn0	O	PCIe TX/RX Differential signals defined by the PCIe 3.1/4.0 specification. The Tx/Rx are defined on module perspective. PCIe Tx net shall connect to the Rx on the host side.		Required
	PETp1/PETn1	O			
	PETp2/PETn2	O			
	PETp3/PETn3	O			
	PERp0/PERn0	I			
	PERp1/PERn1	I			
	PERp2/PERn2	I			
	PERp3/PERn3	I			
	REFCLKp/REFCLKn	I	PCIe Reference Clock signals (100 MHz) defined by the PCIe 3.0/4.0 specification		Required
	PERST#	I	PE-Reset is a functional reset to the card as defined by the PCI Express CEM Rev3.0	3.3V	Required
	CLKREQ#	I/O	Clock Request is a reference clock request signal as defined by the PCIe Mini CEM	3.3V	Optional. The platform will

			specification; Open Drain with pull up on Platform; Active Low; Also used by L1 PM Substates.		leave this pin open.
	PEWAKE#	I/O	Vendor do not need to support this feature	3.3V	Optional. The platform will leave this pin open.
Specific Signals	Reserved for MFG DATA		Manufacturing Data line.		Optional. The platform will leave this pin open.
	Reserved for MFG CLOCK		Manufacturing Clock line.		
	LED1# (O)	O	LED pin	3.3V	Optional. The platform will leave this pin open.
	ALERT#	O	Alert notification to master; Open Drain with pull up on Platform; Active Low.	1.8V	Required. Refer to Sec 5.3 for more details.
	SMB_CLK	I/O	SMBus clock; Open Drain with pull up on Platform, no pull up on module, slave on module	1.8V	Required
	SMB_DATA	I/O	SMBus DATA; Open Drain with pull up on Platform, no pull up on module, slave on module	1.8V	Required
USB	USB_D+	I/O	USB 2.0 bus reserved for future application.	N/A	NC in module, The platform will this pin open.
	USB_D-	I/O	USB 2.0 bus reserved for future application, not required here.	N/A	NC in module, The platform will this pin open.
UART	Reserved_UART_RX	I	UART Receiver Pin. It shall connect to the Tx pin on the host side.	1.8V	Required. Please refer section 5.3 for details.
	Reserved_UART_TX	O	UART Transmitter Pin. It shall connect to the Rx pin on the host side.	1.8V	Required. Please refer section 5.3 for details.

JTAG	TDI	I	Refer to JTAG Specification (IEEE 1149.1), Test Access Port and Boundary Scan Architecture for definition. The definition is also based on module perspective.	1.8V	Required. Please refer section 5.3 for details.
	TDO	O		1.8V	
	TCK	I		1.8V	
	TMS	I		1.8V	
	TRST	I		1.8V	
Reserved New IOs	PWRDIS	I	Reserved for power disable pin. High: disable power on module. This pin shall be NC on module.	3.3V	The platform does not support these features.
	PLN#	I	Reserved for Power Loss notification. NC in module.	3.3V	
	PLA_S3#	O	Reserved for Power loss Assert. NC in module.	3.3V	
	VIO_CFG_GND	O	Reserved for IO configure pin. Connected to ground in module.	0V	

5.2 Power Specification

5.2.1 Operating (steady state) Conditions

The M.2 module utilizes a single regulated power rail of 3.3 V provided by the platform. The power requirement is called out as below:

Table 3. Operation Mode Rating

Norminal supply Voltage	3.3V
Supply Voltage Tolerance	+/-5%
ASIC Junction Temperature	7% lower than the lowest throttle temperature specified at the Max TDP operating point (e.g. if the throttle temp is 70C then the operating temp would be 65C)

The module shall support Low Performance and Target performance levels with Higher and Highest performance mode as an optional level as listed in Table 4. This is a static power level, it is not changed during run-time through software, and defines the maximum sustained power drawn from the system during normal operation. The module components should be designed to support the highest power level that the module can be used in.

Table 4. Module TDP table

	Low Perf	Target Perf	Higher Perf	Highest Perf
Module TDP*	9W	10W	12W	14.85W
Module Absolute Peak Power (up to 20us duration)	15W	18W	20W	24W

Module Performance (assuming DRAM ECC enabled)	N/A	100% of the performance target	>= 110%	>= 120%

*TDP: Thermal Design Power is the sustained average power that the module dissipates while under any work load including synthetic and application work load, 50C local ambient temperature, and the ASIC junction temperature defined above.

5.2.2 Peak (instantaneous) conditions

Transient/Instantaneous power spikes allowed in operation are defined by the absolute peak power specification. Absolute peak power can vary based on the duration time of the peak power transient load step. Fig.8 represents the curve that defines the peak power vs. load step duration time based on the GPv2 card design and Yosemite V2 platform peak power budget. Peak power pulses shorter than 5us shall be supported by the capacitors on the module. These are not specified for a power virus condition, but for peak instantaneous power observed while running application workloads.

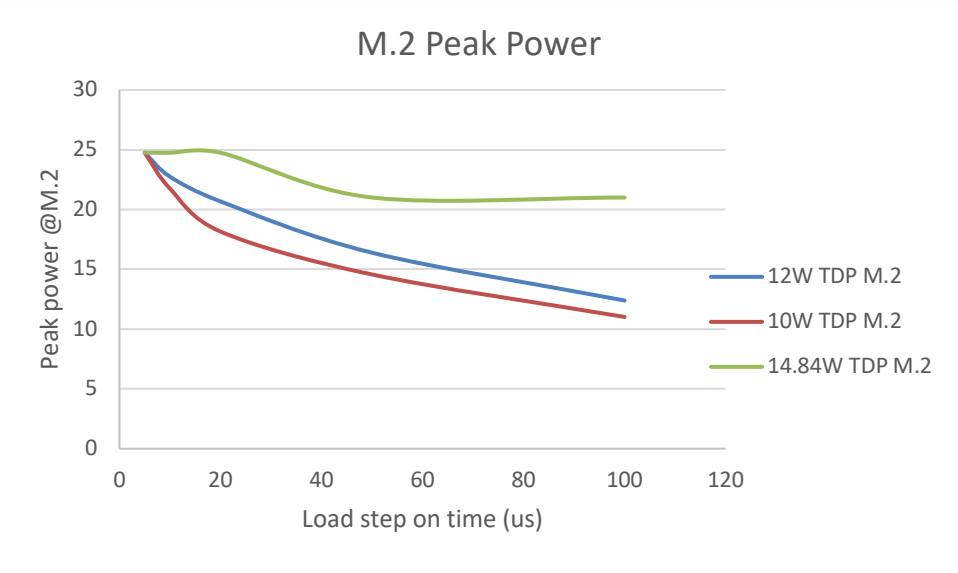

Figure 8.1 M.2 Peak power specification details across load step

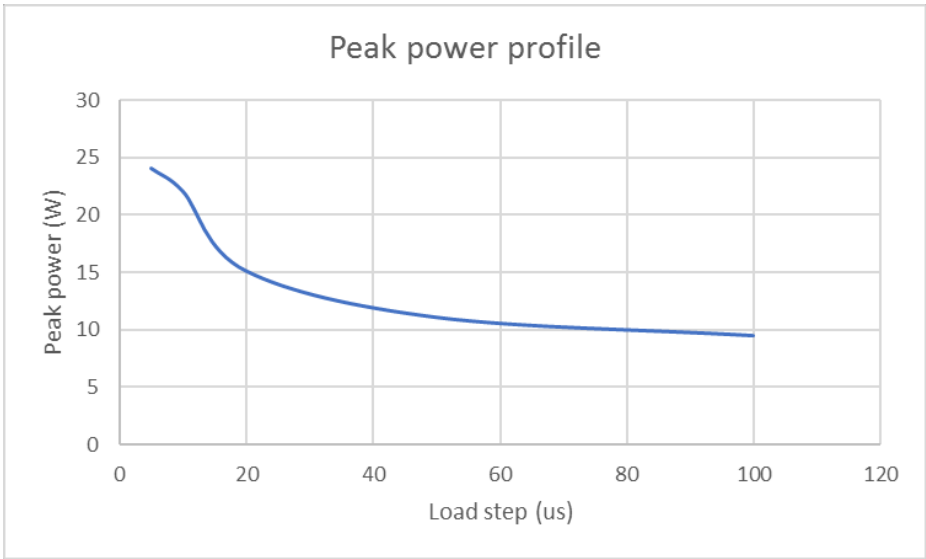

Figure 8.2 M.2 Peak power specification details for low power applications

5.2.1 Input Capacitance and undervoltage specification

Input capacitance on the 3.3V connectors should be limited to less than 1mF per M.2 module. Device UVLO (minimum voltage for M.2 VR to turn on) should be set to greater than 2.8V to ensure power on ramp time meets PCIE ramp spec and does not stress load switch in addition to inrush. Power tree to derive the 3.3V from 12V is described in Figure 9. Power sensors used here are capable of sensing only sustained load (>100ms sampling after averaging).

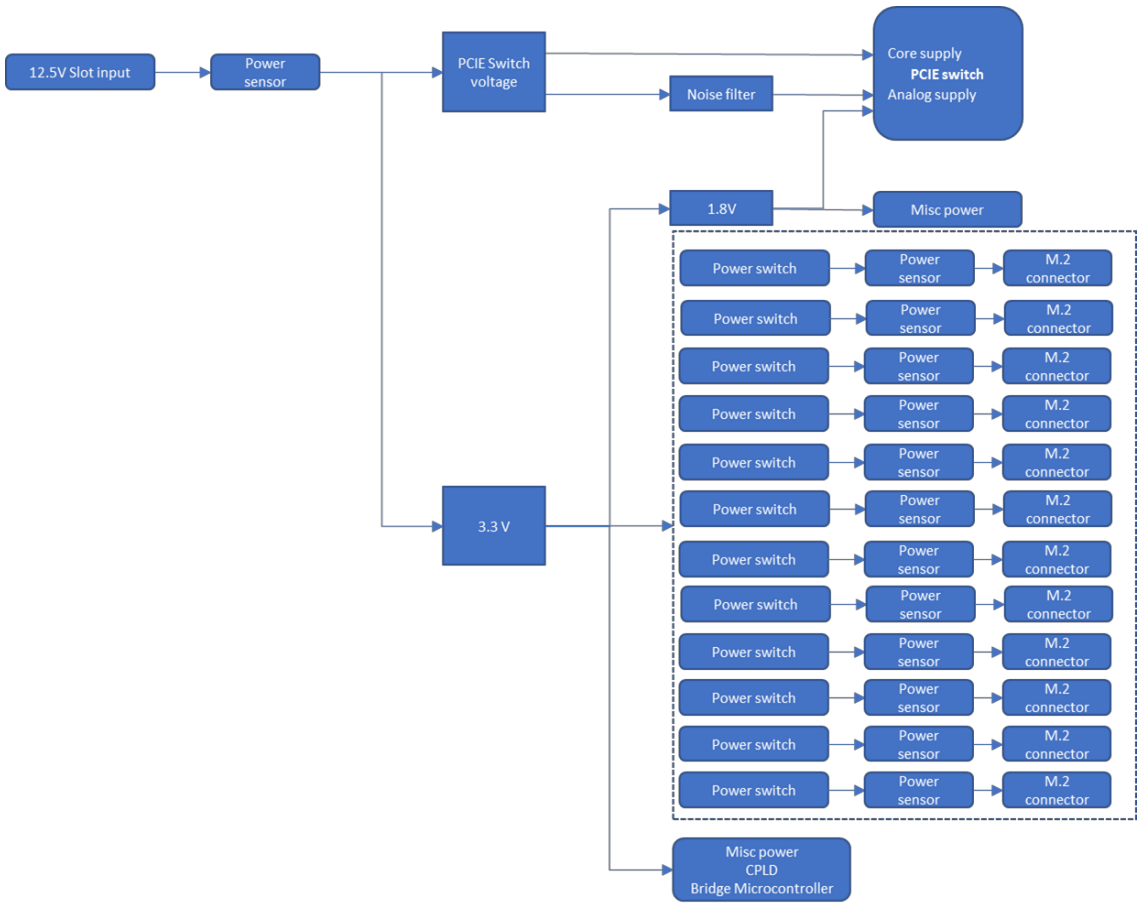

Figure 9. M.2 Power supply configuration on GPv2 carrier card

5.3 IO spec

5.3.1 ALERT#

This pin shall be used as a hardware interrupt for module failures. The ALERT# pin is connected individually to the Bridge IC as shown in Fig. 4. The ALERT# pin shall be triggered by the module health byte as defined in the sideband memory map. The ALERT# pin shall be released after this module health byte is read.

5.3.2 SMBus

The module shall be a SMBus/I2C slave to the host with 7-bit address 0x6A which is defined based on NMVe-MI simplified command interface. FRU information shall be stored in either an EEPROM on the PCB or in on-chip memory in the ASIC, at 8 bit address 0xA6. This bus shall support 100kHz, 400kHz and 1MHz mode. By default, it shall run at the highest speed among the three frequencies mentioned here.

5.3.3 UART

The UART is defined at 1.8V signaling levels. The baud rate is 57600. The module shall buffer this interface to avoid current leakage if the ASIC is not powered.

5.3.4 JTAG

The JTAG interface should be compliant to IEEE standard 1149.1. The module shall buffer this interface to avoid current leakage if the ASIC is not powered.

5.4 PCIe Description

5.4.1 Physical interface

The module PCIe physical interface shall be compliant to PCI-SIG CEM specification 3.0. If the module is capable to run at Gen4 speed, the interface shall be compliant to PCI-SIG CEM Specification 4.0.

The module shall support x4 bifurcation as the default configuration. The module shall support PCIe lane and polarity reversal. Lane reversal is from 0->3 to 3->0.

The module PCIe interface shall support common clock topology with Spread Spectrum Clocking (SSC). SSC's modulation frequency is from 30-33KHz with -0.5%-0% deviation. Separate Reference clock topology support is a preferred but not required.

PCIe routing maximum link loss, including the PCB trace plus the package loss on the module, shall be less than 5dB at 4GHz for PCIe Gen3. The link loss for PCIe Gen4 is defined to be less than 7.5dB at 8GHz. The module trace impedance shall be 85ohm differential impedance for the PCIe lanes.

5.4.2 PCIe power-up timing

The power-up timing of PCIe functions shall follow CEM specification 3.0. Here is the drawing copying from CEM spec:

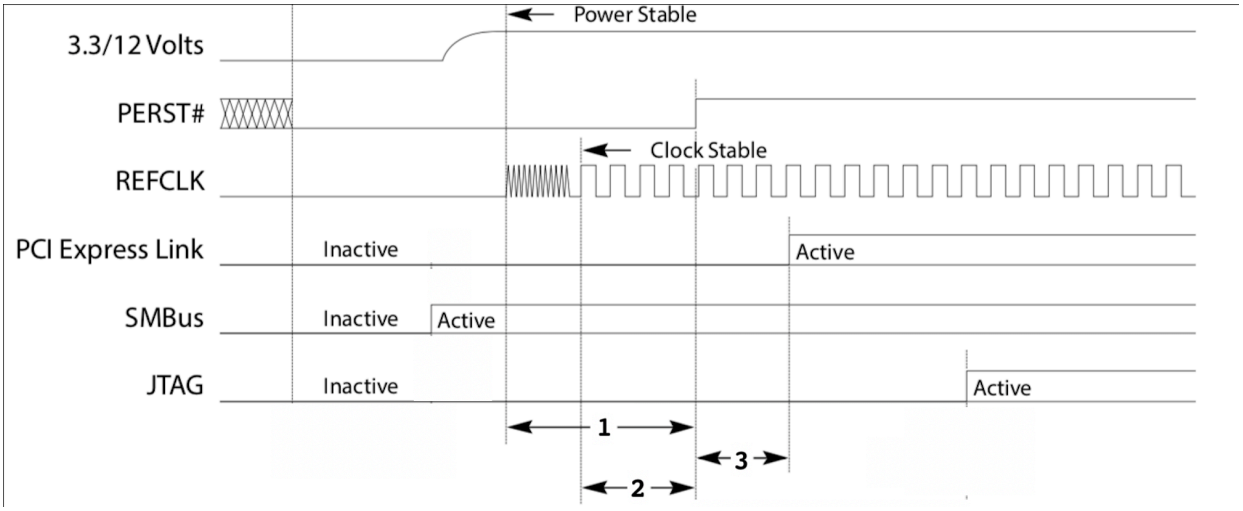

Figure 10: PCIE Power Up Timing

Timing Notes:

1- Minimum time from power rails within specified tolerance to PERST# inactive (T_{PVPERL}): T_{PVPERL} >100ms

2- Minimum clock valid to PERST# inactive($T_{PERST-CLK}$): $T_{PERST-CLK}$ > 100us

3- Minimum PERST# inactive to PCI Express link out of electrical idle:

   a. Link LTSSM must enter detect state within 20ms from PERST# being deasserted. This implies that any PHY settings need to be applied to the PCIe PHY before this timing requirement has been exceeded. Caution should be taken here as this may require a boot ROM to execute and/or information to be loaded from SPI flash as well as a reset sequence applied to the logic for it to take effect.

   b. Link must be ready for configuration request within 120ms from PERST# being de-asserted.

5.4.3 Reset mechanisms

The module shall support the following reset mechanisms when the module POWER GOOD indicates that power supplies are stable on the module.

Out of band:

Cold Reset: Signalled by module power good transition from low to high, followed by PCIe PERST# transition from low to high (refer Fig. 11). This is the cold boot scenario for the module and a fundamental reset as defined by the PCIe specification.

Warm Reset: This is signaled by a transition in PERST# without any transition on module POWER GOOD. This should be interpreted by the ASIC as a warm reset where the PCIe link stays active.

In Band:

PCIe hot reset : This is signaled in band as per the PCIe specification and the device should respond as per the specification.

PCIe Function level reset : This reset is also signaled in-band and the device should meet all the requirements of the PCIe specification.

In addition, as this module will be used in lights out datacenters, we need to have the ability to monitor device health and status using the out of band interface. To meet this requirement the module needs to be able to respond to reads over the OOB interface (SMBUS) while any of the resets are asserted (except COLD RESET). The accelerator shall respond in one of two ways:

1. Return valid data
2. Signal device not ready over the SMBus/i2c protocol

5.4.4 PCIe Configuration

The module shall be configured as a PCIe end-point device. Additionally, the ASIC PCIe controller shall support the following:

1. A Max Payload Size (MPS) of >= 256 Bytes
2. A Max Read Request Size (MRRS) of >= 512 Bytes
3. At least one BAR shall be pre-fetchable, 64bits, and configurable to be at least 1GB in size
4. The DMA engine shall be capable of saturating at least the PCIe gen4 x4 connection (if gen4 capable) or the PCIe gen3 x4 connection (if gen3 capable).

5. The DMA engine shall be capable of mapping to all of host memory and the ability to map the majority of the memory on the module with certain memory regions mapped out due to security concerns.
6. The DMA engine shall support a link latency of >= 1us.
7. The DMA engine shall support the ability for software/firmware to enable/disable PCIe MSI/MSI-X interrupts per DMA command and programmatically map the interrupts to either the host or internal CPU cores so that it is possible to chain multiple PCIe commands together.

5.5 FRU spec

The supplier's FRU is stored in an external EEPROM, or in a memory area within the ASIC that can be accessed from sideband I2C line at address 0xA6. The FRU format should follow *IPMI Platform Management FRU Information Storage Definition 1.0, Version 1.2*. The FRU template is listed in table 5. The supplier should provide the detailed FRU information for review before the module builds.

Table 5. FRU Required Fields

Organization	String
Board Info Area	
Language Code	19h (English)
Board Mfg Date	[Generate build time]
Board Mfg	Defined by vendor
Board Product	Project Code Name
Board Serial Number	Defined by vendor
Board Part Number	Defined by vendor
Fru File ID	Defined by vendor
Custom Field 2	Accelerator M.2
Product Info Area	
Language Code	19h (English)
Product Manufacturer	Defined by vendor
Product Name	Defined by vendor
Part/Model Number	Defined by vendor
Product Version	Defined by vendor
Product Serial Number	Defined by vendor
Product Asset Tag	Defined by vendor
Product Build	EVT (or DVT, PVT)

6 PCB Specification

Refer to PCI Express M.2 Specification Revision 1.1 for the PCB outline mechanical specification.

HDI type PCB manufacturing is expected here to support high density routing and high density BGA package of ASICs. Vendor could use OSP or ENIG surface finishing on PCB except the golden finger area.

The supplier shall follow the requirements below for PCB soldermask color:

EVT-Red

DVT-Yellow

PVT/MP-Blue

7 Thermal and Heatsink

7.1 Thermal Design Guidelines

To improve thermal efficiency, the module shall be fully enclosed by a metal case with a module-level heat sink on the ASIC side. Heat sink dimensions and the associated thermal design requirements need to comply with platform that takes the module. Both the module and the heat sink solution for the module shall be provided by the module supplier. The first platform to use the M.2 accelerator module is Glacier Point V2 (GPV2) carrier card in Yosemite V2 platform (YV2). The information below is intended to serve as a reference for this initial use case only.

7.2 Integrated Heat Sink Requirements for M.2 in GPv2

This section specifies the dimensions for the integrated heat sink solution for the M.2. A reference design is shown in Fig. 11. The supplier is encouraged to use their own module design which meets the mechanical dimension requirements. The latch actuation must be 1.0±0.15 kfg when the latch is fully pressed with 2.5mm travel distance.

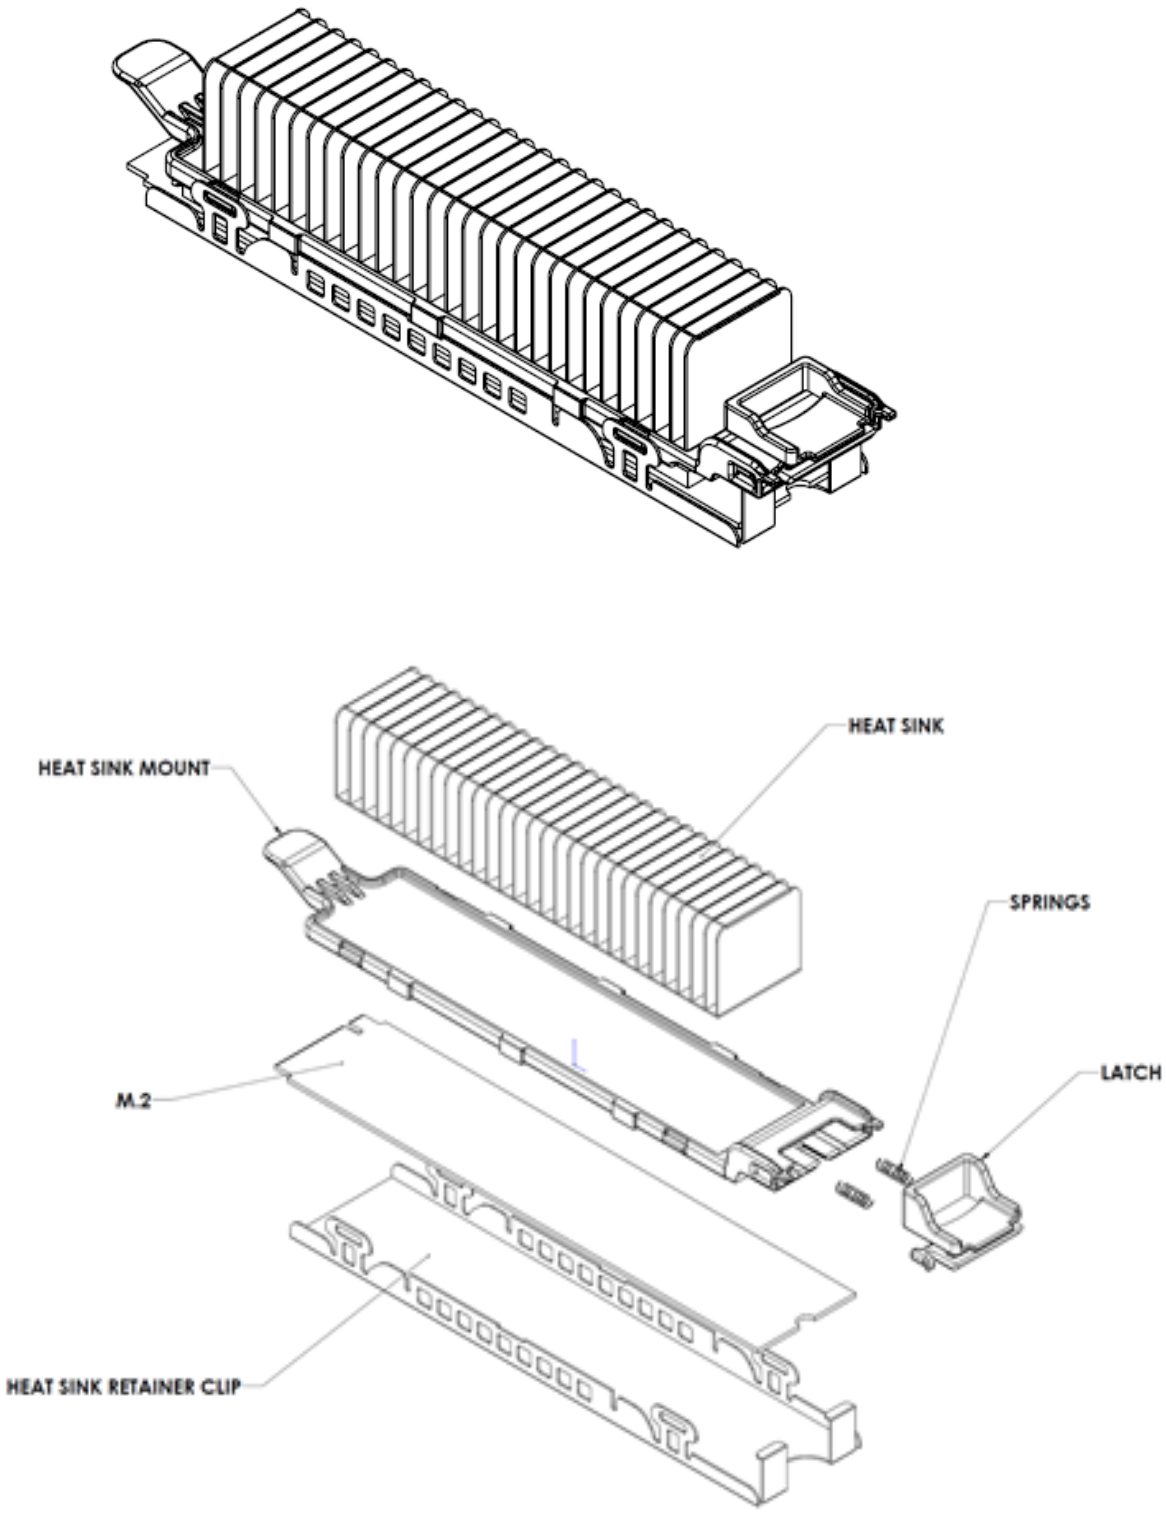

Figure 11. Reference Design for M.2 Integrated heatsink

The heat sink's dimensional requirements for the M.2 accelerator module are listed as below:

- Nominal height from PCB top surface to heat sink top is 21.4mm, which consists of height of heat sink, TIM (thermal interface material) and SMT components on top side.
- Nominal height from PCB bottom surface to bottom case is 2.4mm, which consists of thickness of the metal case, TIM and SMT components on bottom side (if any).
- To provide easy access to connector side and platform integration, the heat sink base of the module shall be die-cast and follow the requirements on the latch design.
- Nominal width of the integrated heatsink shall be kept at 23.4mm for a M.2 module.

The heatsink should follow below tolerance requirements to allow for easy installation into the chassis as shown in Fig. 12.

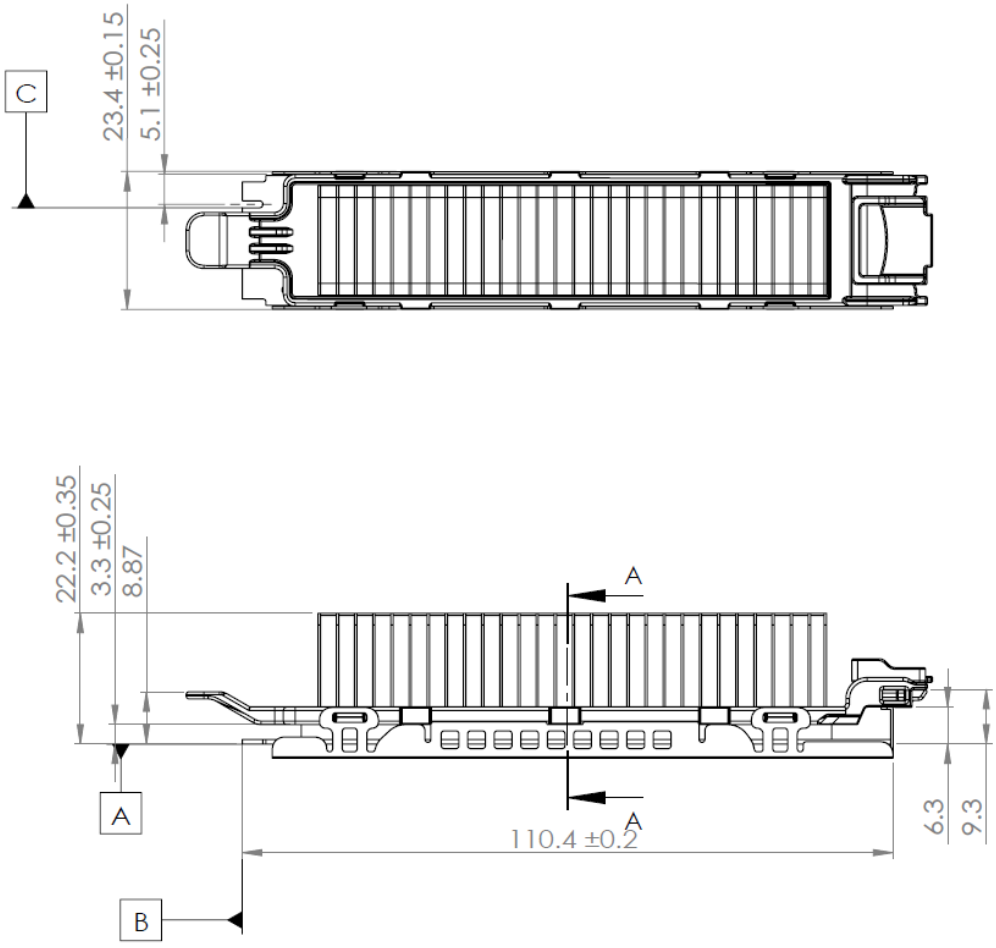

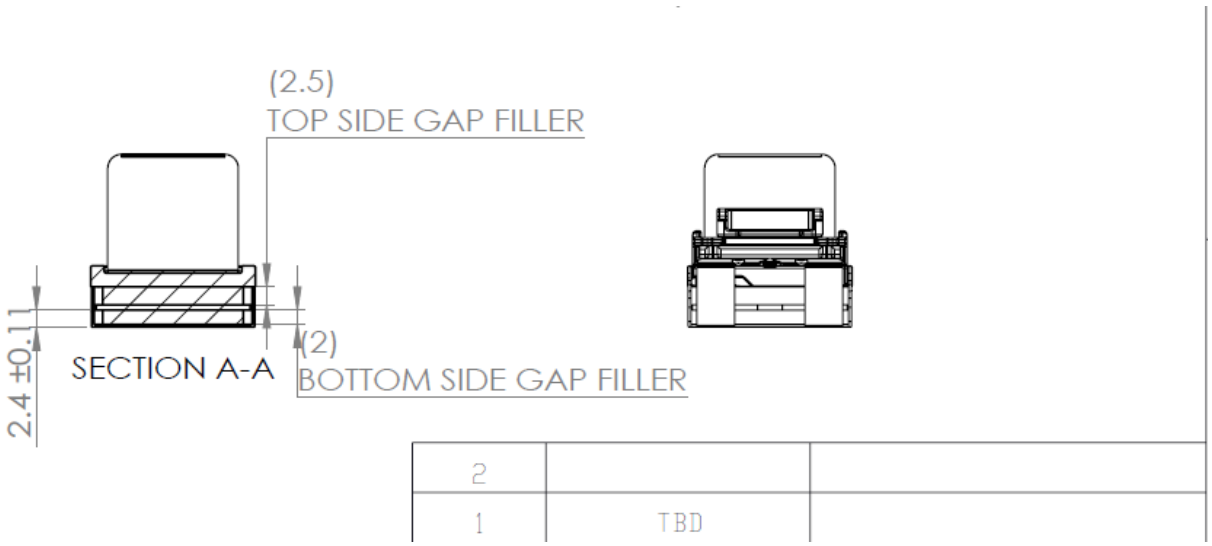

Figure 12. Integrated heatsink dimensions

7.3 Thermal Requirements for M.2 Acceleration Module

The module shall meet the thermal requirement that is defined by the platform. Specific requirements such as airflow and approach air temperature will be platform specific.

All the temperature values reported by module shall hold ±2°C accuracy. ±1°C accuracy is recommended. The ASIC's module-to-module temperature reporting variation shall be ±2°C with ±1°C recommended. Two different modules shall not report temperature greater than 4 °C apart under the same environmental conditions, slot location, and workload.

8 Quality and Reliability:

Module and ASIC are expected to demonstrate RDT, Compliance, Robustness, Environmental specifications, and Component qualification. A mix of industry product quality standards (JESD, ASTM, EIA, ISTA etc.) along with specific product needs that go above and beyond industry standards are required to be completed and demonstrated as part of product quality reliability, and performance.

8.1 RDT

8.1.1 Quality RDT

Hardware component validation, firmware functionality check, as well as product reliability are the key aspects of Reliability Demonstration Test (RDT). The module is expected to demonstrate MTBF (Mean Time Between Failure) of minimum 2.5 million hours in order to assess product early-life quality and potential failure modes. The supplier shall provide the calculated as well as demonstrated MTBF estimates which include sample size, allowable functional failures (1), stress profile (per JESD218A, JESD219A workloads and temperature conditions) at 60%, and 90% confidence level Weibull modeling. In-lieu of calculated MTBF availability, component and module level FIT rates are acceptable as well. Planned and Un-planned power loss scenarios are to be accounted as part of MTBF demonstration.

8.1.2 Performance RDT

Module should run work load that can stress the ASIC, DRAM, and PCIE interface. Module vendor would be required to provide the work load for this purpose. Associated test conditions, performance criteria, and duration for this test are TBD.

8.1.3 Environmental, and Compliance Specifications

Environmental, compliance specifications, and product robustness requirements are listed below so as to ensure that the module/FRU (with Integrated Heat Sink) level requirements are met, and the product functions as expected with no allowable failures post testing.

Table 6. Q&R testing requirement

Module Stress Test	Test Criteria	Standards (As applicable)	Sample Size
Operational Vibration	2.17 G_{rms}, 5-700-5 Hz, all three axes	EIA-364-28	22
Non-Operational Vibration	3.13 G_{rms}, 5-1500-5 Hz, all three axes	EIA-364-28	22
Non-Operational Shock	1250G, 0.5ms, 6 drops, all three axes	EIA-364-27	22
Insertion	300 cycles (plating and power on check every 50 cycles)	EIA-364-09	5
EMC Emission and Immunity		CISPR 22/24, EN55022:2010 +AC:2011, ENTT032:2012 +AC:2013, EN55024:2010 EN 6100-3-2:2014 EN 6100-3-3: 2013 Class B	6
Electrostatic Discharge	± 4kV Contact Discharge ± 8kV Air Discharge	EN55024:2010	6

Module FRU Stress Test (with IHS)	Test Criteria	Standards (Applicable)	Sample Size
Operational Vibration	0.5 G_{rms}, 5-500-5 Hz, all three axes	EIA 364-28	22

Non-Operational Vibration	1.5 G_{rms}, 5-500-5 Hz, all three axes	EIA 364-28	22
Operational Shock	6G, 0.5ms, 6 drops, all three axes	EIA 364-27	22
Non-Operational Shock	70G, 0.5ms, 6 drops, all three axes	MIL-510	22
Package Vibration	1.146 G_{rms}, 2-200-2 Hz, all three axes	ISTA 3E 06-06	10 per tray, 4 trays
Package Drop	8-inch drop	ISTA 3E 06-06	10 per tray, 4 trays
Package Compression	Maximum compression loading on a bulk pack	ASTM D 642-94	1
Thermal Shock **(Non-Operational)**	-40°C to 85°C, 500 cycles (1 Cycle = 5°C to -40°C at ramp 5°C/min, dwell at -40°C for 30 min, -40°C to 85°C at 5°C/min ramp, dwell at 85°C for 30 min, and ramp down to 5°C)	EIA 364-32	22
High Temperature Humidity (Operational)	50°C(local ambient temperature at the module), 90% RH, 500 hours	EIA 364-31	22
Temperature/Voltage Characterization (Operational)	5°C to 50°C (local ambient temperature at the module), Vcc ±x% (per spec), 500 hours		22
Operational Altitude	0 ft to 10,000 ft		12
Non-Operational Altitude	0 ft to 30,000 ft		12
Power Cycle (AC, DC, Reset)	500 cycles each		22

Note: The preference is to make use of at least 3 units from the sample size listed above to be subjected to waterfall model reliability testing (using select few tests from above). Test sequence for this implementation is TBD

8.2 Compliance

North America

- **FCC:** Verification tests only per FCC Part 15 standard. No FCC certification is needed
- **UL:** RU mark is preferred

EU

- **CE mark**: Add CE mark on the accelerator module
- **EMC Directive**: 2014/30/EU- Test partially as applicable
- **ROHS Directive**: 2011/65/EU
- **WEEE mark:** WEEE mark on the accelerator module

APAC

- No specific certification is needed

9 Prescribed Materials

9.1 Disallowed Components

The following components shall not be used in the design of PCB board:

- Components disallowed by the European Union's Restriction of Hazardous Substances Directive **RoHS 2 Directive (2011/65/EU)**
- Trimmers and/or potentiometers
- Dip switches

9.2 Capacitors and Inductors

The following limitations apply to the use of capacitors:

- Only aluminum organic polymer capacitors made by high-quality manufacturers are used; they must be rated 105°C.
- All capacitors have a predicted life of at least 50,000 hours at 45°C inlet air temperature, under the worst conditions.
- Tantalum capacitors using manganese dioxide cathodes are forbidden.
- SMT ceramic capacitors with case size > 1206 are forbidden (size 1206 are still allowed when installed far from the PCB edge and with a correct orientation that minimizes the risk of cracking).
- Ceramic material for SMT capacitors must be X5R or better (COG or NP0 type are used in critical portions of the design). Only SMT inductors may be used. The use of through-hole inductors is disallowed.

9.3 Component De-rating

For inductors, capacitors, and FETs, de-rating analysis is based on at least 20% de-rating.

10 Labels and Markings

This specification describes label requirements for M.2 accelerators.

10.1 Data required

- Manufacturer name
- Country of Origin
- Date code of manufacture, includes year & work week
- Product number = same number used to order product from supplier
- Serial number: unique to each product
- Firmware revision
- Hardware revision
- Capacity of card (GB) – TBD: total data space (system + user), total user space, or user space less OP. (This request is for storage product only)
- PCB Vendor name

We might need to apply TIM (gap pad) material on top of high heat dissipating components such as ASIC, NAND flash and controller on either side of the card. Any standard product labels applied on those parts will be obscured by the TIM, so the supplier should install additional labels that contain the Serial Number and the Product Number (same PN used for ordering and stored in the SMART data table. The additional labels can use 2D barcodes to save space. 2D bar codes should have human readable text placed in the margin. 2D barcodes should not have any spaces, but dashes are acceptable.

In another case we may have integrated heat sink that encloses module. In this case the label should be attached on bottom of integrated heat sink.

10.2 Data format

- Human-readable. Font size: 10 or larger. Some data on 2D labels can be size 6
- Barcode. 1D and 2D acceptable. Minimum feature (line width): 10 mils
- Electronically readable, e.g. SMART data table
- Example readers: Motorola/Zebra readers
    - Motorola CS4070
    - Motorola Symbol DS3578-SR
    - Zebra DS3678-DP

Table 7. Label Requirements

	Spec		
Data	Human Readable	Barcode	Electronic
Product name	X	X	X

Capacity ()			
Serial Number (Human Readable)	X		X
Serial Number (Barcode)		X	
Sub-assembly No. (Human Readable)			
Sub-assembly number (Barcode)		X	
PCBA Number			
LBA			
Country of Manufacture	X		
Model String	X	Highly Wanted	X
Warranty Disclaimer			
WWN Worldwide Name (human Readable)			
WWN Worldwide Name (Barcode)			
Firmware Version	X		X
Canadian String			
Manufacturer Name	X	X	X
PCB Vendor		X	X
Date code		X	X
PSID Human readable			
PSID Barcode			
Production Date Code	X		X

10.3 Agency Compliance Marks

Module supplier should ensure its products comply with all applicable certificate(s) or verification among the following requirements.

EMC/Safety

- NRTL component level certification
- CE mark based on Directive 2014/35/EC and 2014/30/EU
- FCC verification
- VCCI
- Korea KCC certification
- Taiwan BSMI

Environment regulations

- ROHS - Must be free from hazardous substances prohibited by the RoHS Directives of the EU (European Union)
- China RoHS
- Taiwan RoHS
- WEEE mark
- Halogen Free

11 Revision History

Ver	Description	Author	Date
0.1	Initial submission	Facebook	02/26/2019