Intel® Data Center GPU Max 1100 Accelerator

Datasheet

April 2024

Revision 1.0

Document Number: 817799

Intel technologies may require enabled hardware, software or service activation.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site.

Your costs and results may vary.

"Conflict-free" refers to products, suppliers, supply chains, smelters, and refiners that, based on our due diligence, do not contain or source tantalum, tin, tungsten or gold (referred to as "conflict minerals" by the U.S. Securities and Exchange Commission) that directly or indirectly finance or benefit armed groups in the Democratic Republic of the Congo or adjoining countries.

All product plans and roadmaps are subject to change without notice.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Altering clock frequency or voltage may void any product warranties and reduce stability, security, performance, and life of the processor and other components. Check with system and component manufacturers for details.

Results have been estimated or simulated.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visiting the Intel Resource and Document Center.

Contents

Figures

Tables

Revision History

Revision Number	Description	Date
1.0	• Initial public document release.	April 2024

1 Overview

This document provides architectural, thermal, mechanical, and electrical specifications for the Intel® Data Center GPU Max 1100 Add-In Card (AIC). See Figure 1-1 for information on the Intel Data Center GPU Max 1100 accelerator.

Figure 1-1. Intel® Data Center GPU Max 1100 Accelerator Specification

Package Configuration	
SKU	Intel® Data Center GPU Max 1100 300W Add-in Card 1 Stack 56 Xe Cores 448 Vector Engines 448 Matrix Engines
Memory	48GB HBM2e L1: 512KB per Xe Core. 28MB total L2 Cache: 108MB ECC support
Card TDP	300W
Form Factor	 Passive Air-Cooled Double Wide PCIe Card. 10.5" Long
Scalability	Xe Link (Six 53Gb Ports). Supports all-to-all connections between two or four cards via an optional bridge card.
I/O Host Interface	PCIe Gen 5
Compute Numerics	Systolic: Int4, Int8, Bfloat16, FP16, TF32 Non Systolic: Int16, Int32, FP16, FP32, FP64

Notes:

- The content of this document is subject to changes and additions. See the Resource and Documentation Center (RDC) for the latest Datasheet revision.

1.1 References

Specifications and concepts from the following documents and projects are referred to in this document.

Table 1-1. Reference Documents

Document	Document Location
PCI Express Card Electromechanical Specification Revision 5.0	https://pcisig.com/
HBM2e Specification (JESD235D)	https://www.jedec.org/
Open Compute Project (OCP) - Various related specifications	https://www.opencompute.org/
Distributed Management Task Force (DMTF*)* - Open manageability standards	https://www.dmtf.org/

1.2 Terminology

The following initialisms and terms are used in this document:

Table 1-2. Terms and Acronyms

Term	Description
AIC	Add-In Card. Refers to cards designed to plug into a PCI Express* (PCIe*) interface.
AMC	Add-In-Card Management Controller
BMC	Baseboard Management Controller
CEM	Card Electromechanical Specification
DTS	Digital Temperature Sensors
EU	Execution Unit
FRU	Field Replaceable Unit
GPU	Graphics Processing Unit
HBM2e	High Bandwidth Memory 2 evolution (3.6 Gbps)
Intel® OOBM Services Module	Intel® Out-of-Band Management Services Module
Intel® X^e Link	High-Speed coherent unified fabric connecting multi-GPU configurations
MSI	Message Signaled Interrupt
OOB	Out of Band
PCIe*	PCI Express*
PLL	Phase Locked Loop
PUNIT	Power Unit
RAPL	Running Average Power Limit
RAS	Reliability, Availability and Serviceability
SKU	Stock Keeping Unit
SMBus	System Management Bus

Term	Description
SoC	System on Chip
SPI	Serial Peripheral Interface
SVID	Serial Voltage Identification
TBD	To Be Decided
TDC	Thermal Design Current
TDP	Thermal Design Power
VR	Voltage Regulator
Xe Stack	Replaces the name *tile*. Includes cores, memory, RT units, cache, HBM2e controllers, and so forth.

1.3 Related Documentation

Additional product details and specifications are available in the following documents.

Table 1-3. Reference Documents

Document	Document Number
Intel® Data Center GPU Max Series - Management Specification	633513
Intel® Data Center GPU Max Series Specification Update	795271
Intel® Xe Link Bridge Card Datasheet	788941

2 Architecture Overview

2.1 Form Factors

The Intel Data Center GPU Max 1100 and mechanical form factor was designed following the *PCI Express Card Electromechanical Specification* 5.0. However, there are some functional deviations preventing the device from obtaining the PCI Express* (PCIe*) compliance logo, outlined in the *Intel® Data Center GPU Max Series Specification Update*, document number 795271. The Intel Data Center GPU Max 1100 AIC SKU is defined in the following table.

Table 2-1. Intel® Data Center GPU Max Series AIC SKU Information

SKU	Card Size	Card Total Board Power (TBP)	Number of GPUs	Number of Execution Units (EUs)	HBM2e Memory Capacity	PCI Device ID
1100	Full height, three-quarter length, dual slot width	300W	1	448 Matrix and Vector	48 GB	0x0BDA

NOTE: SKU information can be obtained OOB by reading the card's Field Replaceable Unit (FRU) data structure. See the *Intel® Data Center GPU Max Series - Management Specification,* document number 633513, for more information.

2.2 Intel® Xe Link Interconnects

The Intel Data Center GPU Max 1100 can run in one, two, or four card configurations. Two and four-card configurations can use Intel® Xe Link connections for direct all-to-all card-to-card communication. See Section 3.3 for Intel Xe Link topology details.

2.3 Card Block Diagrams

The following figure outlines the major components on the Intel Data Center GPU Max
1100 card.

Figure 2-1. Intel Data Center GPU Max 1100 Card Block Diagram

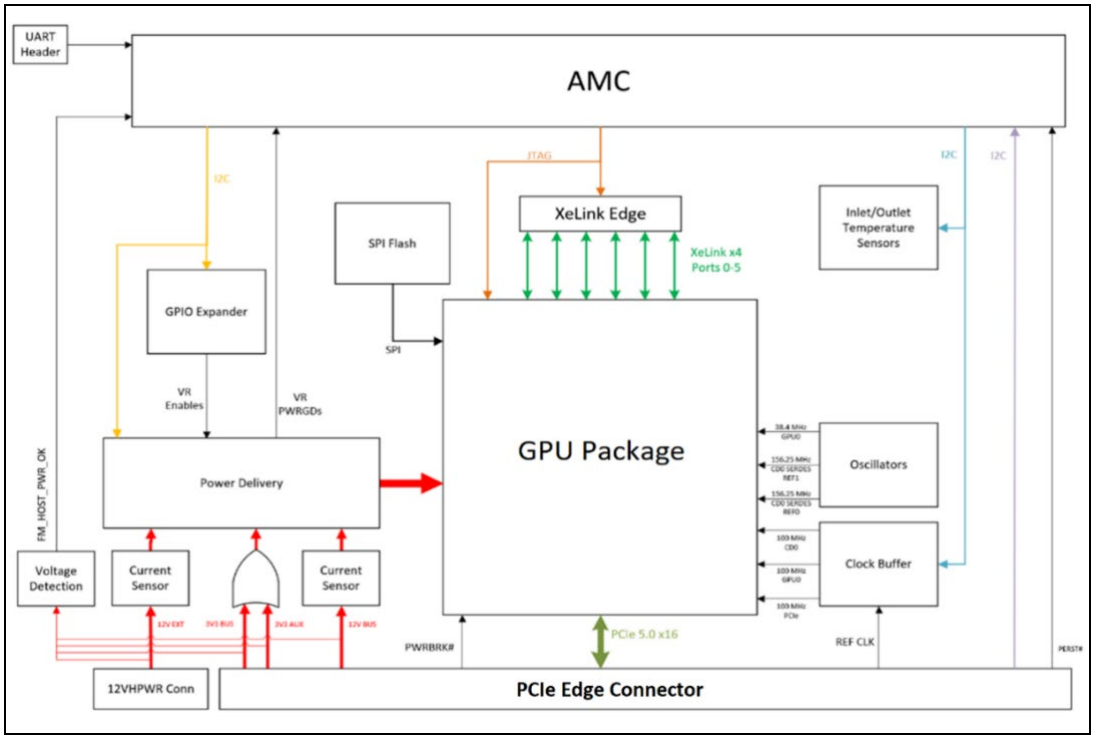

Notes:

- The Intel Xᵉ Link edge connector features six 53 Gbps Serializer/Deserializer (SerDes) ports for all-to-all connections between two or four cards.
- An onboard microcontroller, called the Add-In-Card Management Controller (AMC), controls card power and reset sequencing and provides an SMBus interface for the host to communicate with the card for status, sensor monitoring, FRU access, debug information, and other functions. For details, see the *Intel® Data Center GPU Max Series - Management Specification*, document 633513.
- An optional PWRBRK# signal directly connects from a PCIe edge finger to the GPU package to allow emergency power throttling from the host.
- Temperature sensors enable monitoring and managing card thermals.
- VRs powered from the auxiliary power cable and baseboard PCIe slot.
- Programmable power capping from 1.2x TDP to 2.0x TDP. Default 1.52x.
- CEM 5.1 12V-2x6 H++ auxiliary power input connector.
- PCIe 5.0 (x16 32GT/s) interface to the host platform.
- Three 16 GB HBM2e memory stacks (3.2 GT/s per HBM2e stack) (1.6 TB/s peak BW).

2.4 High-Bandwidth Memory (HBM)

The Intel Data Center GPU Max 1100 features three enabled in-package 16 GB HBM2E stacks. The GPU physically contains four stacks, but one stack is permanently disabled. The disabled stack can be any one of the four and will vary from card to card. HBM sensor readings for the disabled stack will appear in telemetry output, but readings reflect 0 or FF.

2.5 Single Root I/O Virtualization (SR-IOV)

The Intel Data Center GPU Max 1100 supports up to 63 virtual functions per card. For using more than seven virtual functions, Alternative Routing ID (ARI) support is required for functionality. PCI config space Base Address Registers (BARs) for the physical and virtual functions are shown in the following table.

Table 2-2. BAR Sizes, Physical and Virtual

Title	Item
BAR Size (Physical function)	BAR0: 16 MB BAR2: 64 GB† (total HBM size [48 GB]) BAR2 (32 bit): 4 MB
BAR Size (Virtual function)	BAR0/1: 1 GB† (16 MB × 63) BAR2/3: 512 GB (8 GB × 63) (Resizable)

NOTE: BARs are allocated in powers of two. Values marked with "†" are rounded up to the nearest power to reflect the actual address space allocated by BIOS.

2.6 AIC PCI Device IDs

The Intel Data Center GPU Max 1100 has seven PCI device IDs within the GPU. All device IDs are OS visible with or without the i915 driver module loaded. Device 0bdd is the upstream device ID that reflect the card's host link speed and width. When correctly trained, the link should reflect x16 at 32 GT/s. GPU device 0bda will link at x1 2.5 GT/s and reflects the device BAR. See the following figure for the device IDs organized by connection.

Figure 2-2. AIC PCI Device IDs Organized by Connection

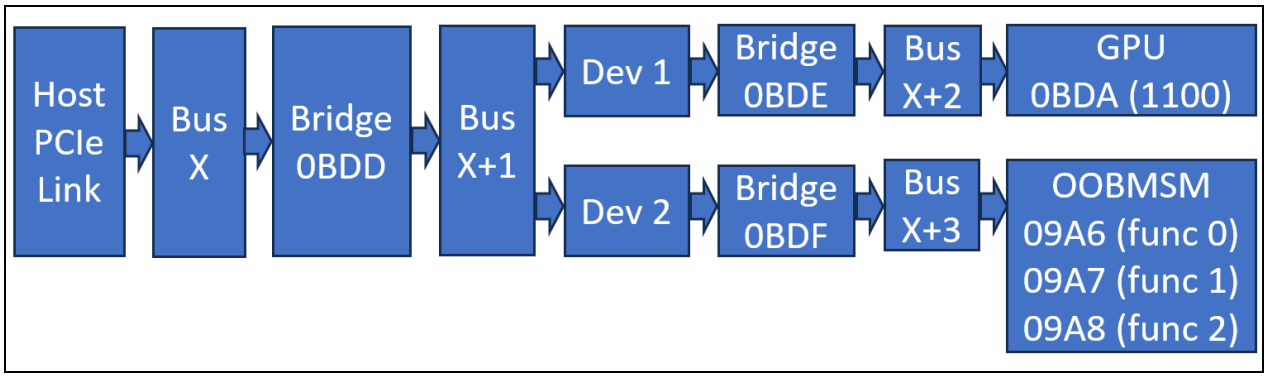

2.7 Supported Message Signaled Interrupt Types

The Intel Data Center GPU Max 1100 supports MSI. MSI-X is not supported.

3 Intel Data Center GPU Max 1100 Card Interfaces

The Intel Data Center GPU Max 1100 card features an Intel Xe Link connector for direct all-to-all card communication and a PCIe 5.0 edge connector for host communication. The card draws power from the PCIe slot and an auxiliary 12V-2x6 H++ power connector. In addition, the card features a UART connector for collecting AMC debug information and performing AMC updates when OOB update is not available. Refer to the appendix at the end of this document for UART details. Connector locations are shown in the following figure.

Figure 3-1. Intel Data Center GPU Max 1100 Interface Locations

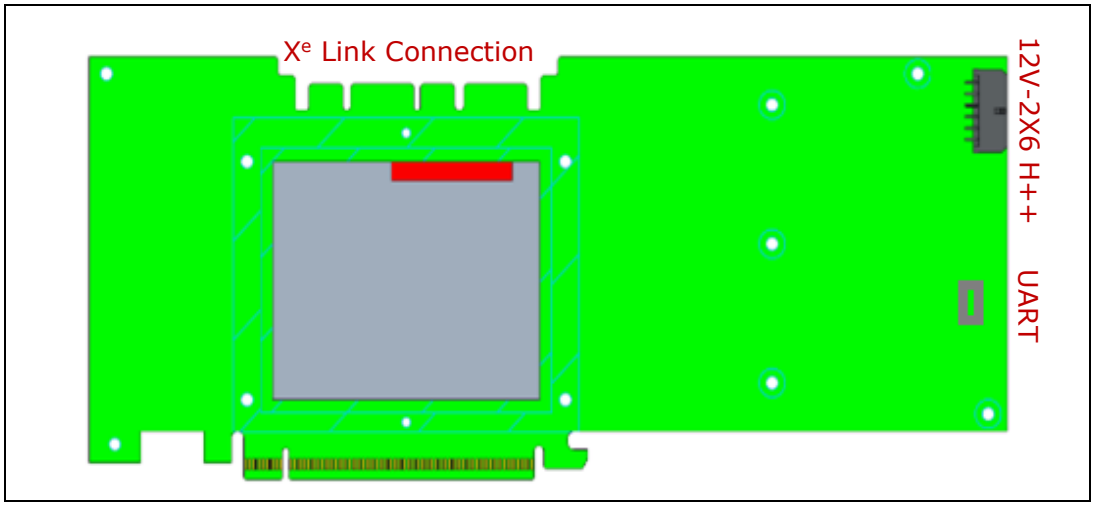

3.1 PCIe Edge Finger Pin List

The following table lists the edge finger pins and their descriptions.

Table 3-1. PCIe Edge Connector Pin Description

Pin	Signal	Description	Pin	Signal	Description
B1	P12V	+12V Power Input	A1	PRSNT1#	AC Termination on Presence pin
B2	P12V	+12V Power Input	A2	P12V	+12V Power Input
B3	P12V	+12V Power Input	A3	P12V	+12V Power Input
B4	GND	Ground	A4	GND	Ground
B5	SMB_DEVICE_PCIE_SCL	System Management Bus (SMBus) Clock to Add-In-Card Management Controller (AMC)	A5	TP_GF_JTAG2	Unused JTAG* TCK
B6	SMB_DEVICE_PCIE_SDA	SMBus Data to/from AMC	A6	TP_GF_JTAG3	Unused JTAG TDI
B7	GND	Ground	A7	TP_GF_JTAG4	Unused JTAG TDO
B8	P3V3	+3.3V Power Input	A8	TP_GF_JTAG5	Unused JTAG TMS
B9	TP_GF_JTAG1	Unused Joint Test Action Group (JTAG) TRST#	A9	P3V3	+3.3V Power Input
B10	P3V3_PCIE_STBY	+3.3V Standby Power Input	A10	P3V3	+3.3V Power Input
B11	IRQ_LVC3_WAKE_N_SLOT5	Unused WAKE#	A11	RST_HOST_PERST_IN_N	PCIe Reset signal from Host
B12	CLKREQ_PCIE_N	Unused CLKREQ#; AC terminated	A12	GND	Ground
B13	GND	Ground	A13	CLK_100M_PCIE_P	PCIe Reference Clock Input
B14	P5E_PVC_PE0_RX_DP_15	Host Tx, ATS SOC Rx Lane 0 P/N	A14	CLK_100M_PCIE_N	
B15	P5E_PVC_PE0_RX_DN_15		A15	GND	Ground
B16	GND	Ground	A16	P5E_PVC_PE0_TX_DP_15	Host Rx, ATS SOC Tx Lane 0 P/N
B17	PRSNT2_PCIE_R_N1	AC Termination on Presence pin	A17	P5E_PVC_PE0_TX_DN_15	
B18	GND	Ground	A18	GND	Ground
B19	P5E_PVC_PE0_RX_DP_14	Host Tx, ATS SOC Rx Lane 1 P/N	A19	GF1_A19	AC terminated
B20	P5E_PVC_PE0_RX_DN_14		A20	GND	Ground
B21	GND	Ground	A21	P5E_PVC_PE0_TX_DP_14	Host Rx, ATS SOC Tx Lane 1 P/N
B22	GND	Ground	A22	P5E_PVC_PE0_TX_DN_14	
B23	P5E_PVC_PE0_RX_DP_13	Host Tx, ATS SOC Rx Lane 2 P/N	A23	GND	Ground
B24	P5E_PVC_PE0_RX_DN_13		A24	GND	Ground
B25	GND	Ground	A25	P5E_PVC_PE0_TX_DP_13	Host Rx, ATS SOC Tx Lane 2 P/N
B26	GND	Ground	A26	P5E_PVC_PE0_TX_DN_13	
B27	P5E_PVC_PE0_RX_DP_12	Host Tx, SOC Rx Lane 3 P/N	A27	GND	Ground
B28	P5E_PVC_PE0_RX_DN_12		A28	GND	Ground
B29	GND	Ground	A29	P5E_PVC_PE0_TX_DP_12	Host Rx, ATS SOC Tx Lane 3 P/N
B30	PCIE_PWRBRK_N	Optional Emergency Power Reduction In	A30	P5E_PVC_PE0_TX_DN_12	
B31	PRSNT2_PCIE_R_N2	AC Termination on Presence pin	A31	GND	Ground
B32	GND	Ground	A32	GF1_A32	AC terminated
B33	P5E_PVC_PE0_RX_DP_11	Host Tx, ATS SOC Rx Lane 4 P/N	A33	GF1_A33	AC terminated
B34	P5E_PVC_PE0_RX_DN_11		A34	GND	Ground
B35	GND	Ground	A35	P5E_PVC_PE0_TX_DP_11	Host Rx, ATS SOC Tx Lane 4 P/N
B36	GND	Ground	A36	P5E_PVC_PE0_TX_DN_11	
B37	P5E_PVC_PE0_RX_DP_10	Host Tx, ATS SOC Rx Lane 5 P/N	A37	GND	Ground
B38	P5E_PVC_PE0_RX_DN_10		A38	GND	Ground
B39	GND	Ground	A39	P5E_PVC_PE0_TX_DP_10	Host Rx, ATS SOC Tx Lane 5 P/N
B40	GND	Ground	A40	P5E_PVC_PE0_TX_DN_10	

Pin	Signal	Description	Pin	Signal	Description
B41	P5E_PVC_PE0_RX_DP_9	Host Tx, ATS SOC Rx Lane 6 P/N	A41	GND	Ground
B42	P5E_PVC_PE0_RX_DN_9		A42	GND	Ground
B43	GND	Ground	A43	P5E_PVC_PE0_TX_DP_9	Host Rx, ATS SOC Tx Lane 6 P/N
B44	GND	Ground	A44	P5E_PVC_PE0_TX_DN_9	
B45	P5E_PVC_PE0_RX_DP_8	Host Tx, ATS SOC Rx Lane 7 P/N	A45	GND	Ground
B46	P5E_PVC_PE0_RX_DN_8		A46	GND	Ground
B47	GND	Ground	A47	P5E_PVC_PE0_TX_DP_8	Host Rx, ATS SOC Tx Lane 7 P/N
B48	PRSNT2_PCIE_R_N3	AC Termination on Presence pin	A48	P5E_PVC_PE0_TX_DN_8	
B49	GND	Ground	A49	GND	Ground
B50	P5E_PVC_PE0_RX_DP_7	Host Tx, ATS SOC Rx Lane 8 P/N	A50	GF1_A50	AC terminated
B51	P5E_PVC_PE0_RX_DN_7		A51	GND	Ground
B52	GND	Ground	A52	P5E_PVC_PE0_TX_DP_7	Host Rx, ATS SOC Tx Lane 8 P/N
B53	GND	Ground	A53	P5E_PVC_PE0_TX_DN_7	
B54	P5E_PVC_PE0_RX_DP_6	Host Tx, ATS SOC Rx Lane 9 P/N	A54	GND	Ground
B55	P5E_PVC_PE0_RX_DN_6		A55	GND	Ground
B56	GND	Ground	A56	P5E_PVC_PE0_TX_DP_6	Host Rx, ATS SOC Tx Lane 9 P/N
B57	GND	Ground	A57	P5E_PVC_PE0_TX_DN_6	
B58	P5E_PVC_PE0_RX_DP_5	Host Tx, ATS SOC Rx Lane 10 P/N	A58	GND	Ground
B59	P5E_PVC_PE0_RX_DN_5		A59	GND	Ground
B60	GND	Ground	A60	P5E_PVC_PE0_TX_DP_5	Host Rx, ATS SOC Tx Lane 10 P/N
B61	GND	Ground	A61	P5E_PVC_PE0_TX_DN_5	
B62	P5E_PVC_PE0_RX_DP_4	Host Tx, ATS SOC Rx Lane 11 P/N	A62	GND	Ground
B63	P5E_PVC_PE0_RX_DN_4		A63	GND	Ground
B64	GND	Ground	A64	P5E_PVC_PE0_TX_DP_4	Host Rx, ATS SOC Tx Lane 11 P/N
B65	GND	Ground	A65	P5E_PVC_PE0_TX_DN_4	
B66	P5E_PVC_PE0_RX_DP_3	Host Tx, ATS SOC Rx Lane 12 P/N	A66	GND	Ground
B67	P5E_PVC_PE0_RX_DN_3		A67	GND	Ground
B68	GND	Ground	A68	P5E_PVC_PE0_TX_DP_3	Host Rx, ATS SOC Tx Lane 12 P/N
B69	GND	Ground	A69	P5E_PVC_PE0_TX_DN_3	
B70	P5E_PVC_PE0_RX_DP_2	Host Tx, ATS SOC Rx Lane 13 P/N	A70	GND	Ground
B71	P5E_PVC_PE0_RX_DN_2		A71	GND	Ground
B72	GND	Ground	A72	P5E_PVC_PE0_TX_DP_2	Host Rx, ATS SOC Tx Lane 13 P/N
B73	GND	Ground	A73	P5E_PVC_PE0_TX_DN_2	
B74	P5E_PVC_PE0_RX_DP_1	Host Tx, ATS SOC Rx Lane 14 P/N	A74	GND	Ground
B75	P5E_PVC_PE0_RX_DN_1		A75	GND	Ground
B76	GND	Ground	A76	P5E_PVC_PE0_TX_DP_1	Host Rx, ATS SOC Tx Lane 14 P/N
B77	GND	Ground	A77	P5E_PVC_PE0_TX_DN_1	
B78	P5E_PVC_PE0_RX_DP_0	Host Tx, ATS SOC Rx Lane 15 P/N	A78	GND	Ground
B79	P5E_PVC_PE0_RX_DN_0		A79	GND	Ground
B80	GND	Ground	A80	P5E_PVC_PE0_TX_DP_0	Host Rx, ATS SOC Tx Lane 15 P/N
B81	PRSNT2_PCIE_R_N4	Connected via 0Ω resistor to pin A1	A81	P5E_PVC_PE0_TX_DN_0	
B82	GF1_B82	AC terminated	A82	GND	Ground

3.2 Intel Data Center GPU Max 1100 Intel Xe Link Connector Options

Intel Xe Link bridge cards connect two or four Intel Data Center GPU Max 1100 cards to enable direct card-to-card all-to-all communication, as detailed in this section.

The bridge cards:

- Sit flush with the PCIe cards' edge after installation and mechanically meet PCIe CEM compliance in all dimensions.
- Support Intel Xe Link data rates up to 53 Gbps.
- Are available for x2 and x4 card configurations.

X2 and X4 Bridge cards provide different amounts of connections between devices as detailed in the following sections, but in summary:

- X2 bridges connect two cards and provide six Intel Xe Link connections between them.
- X4 bridges connect four cards and provide two Intel Xe Link connections between each card.

3.2.1 Intel Data Center GPU Max 1100 Two-Card Topology

Figure 3-2. Intel Xe Link Topology for Two-Card Systems

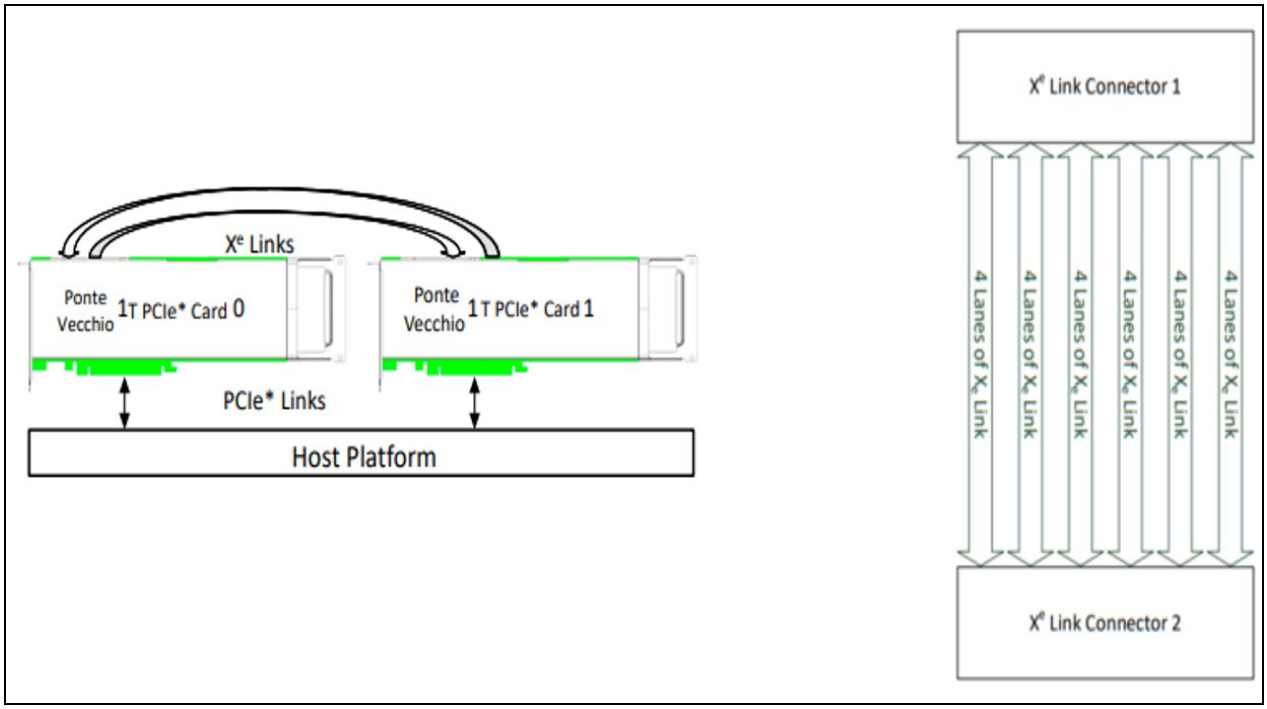

Figure 3-3. Two-Card Configuration Numbering with Intel Xe Link Physical Bridge Installed

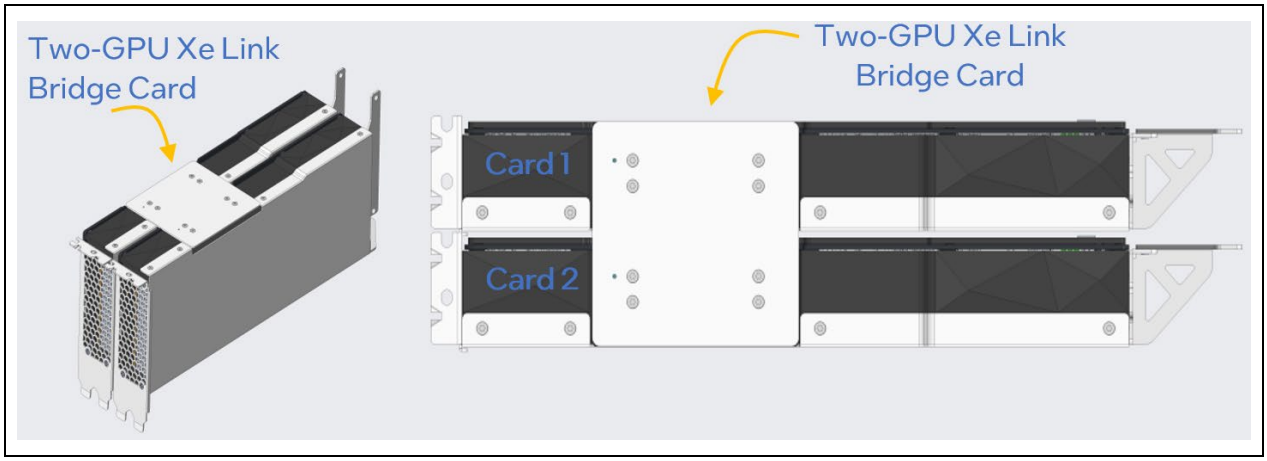

Figure 3-4. Two-Card Intel Xe Link Bridge Connector

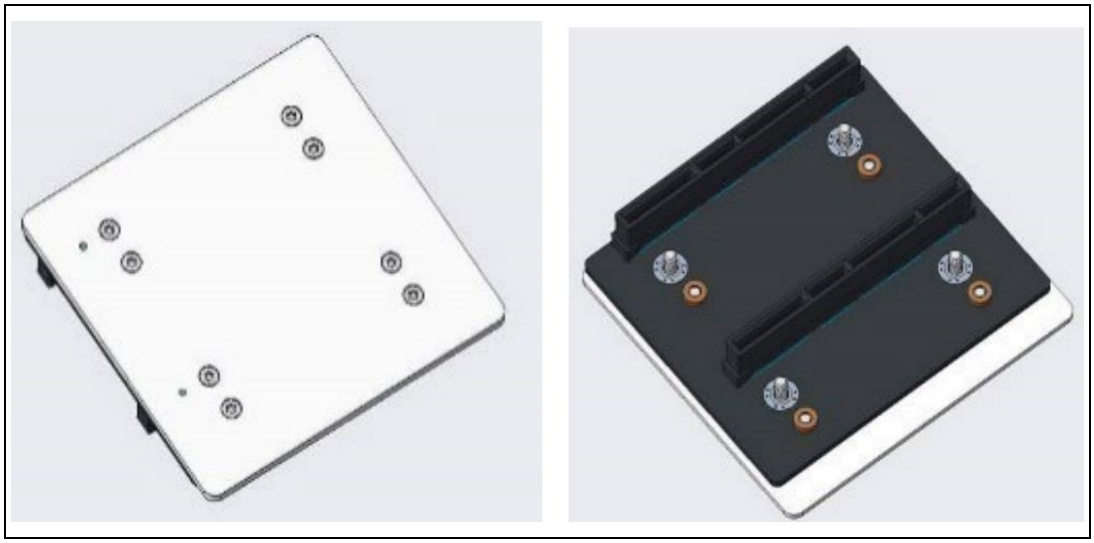

3.2.2 Intel Data Center GPU Max 1100 Four-Card Topology

Figure 3-5. Intel Xe Link Topology for Four-Card System

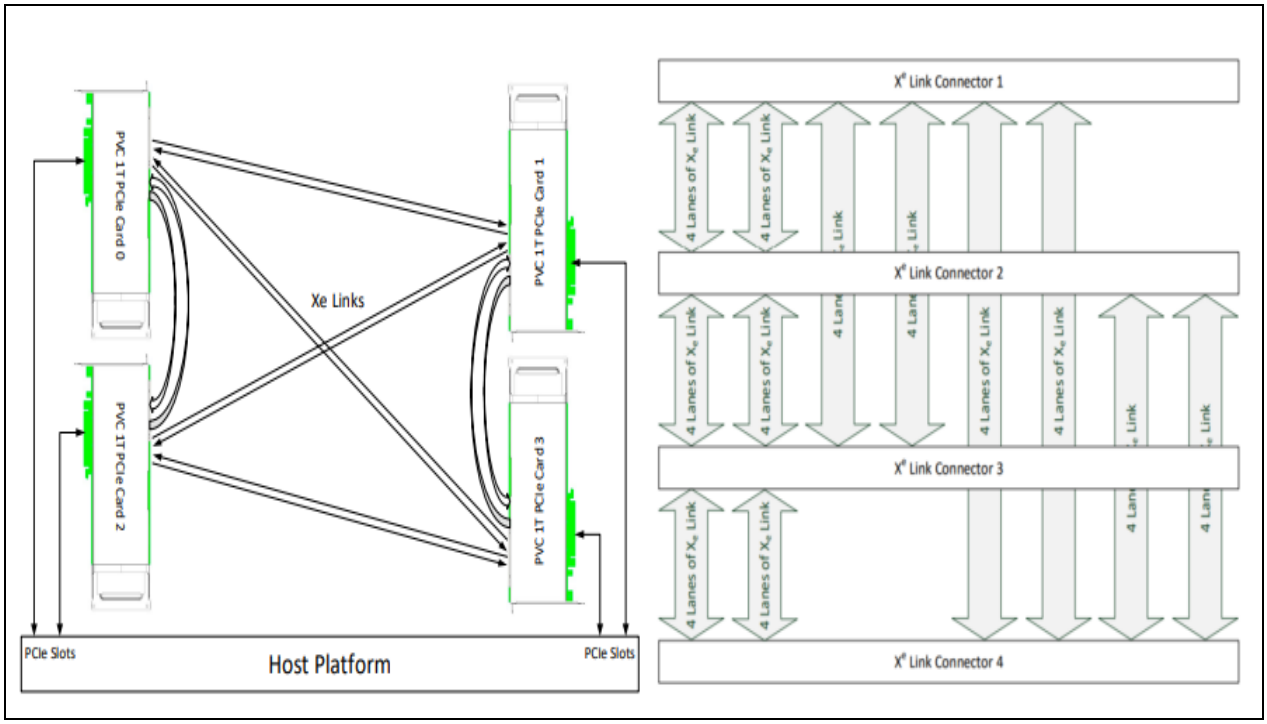

Figure 3-6. Four-Card Configuration Numbering with Intel Xᵉ Link Physical Bridge Installed

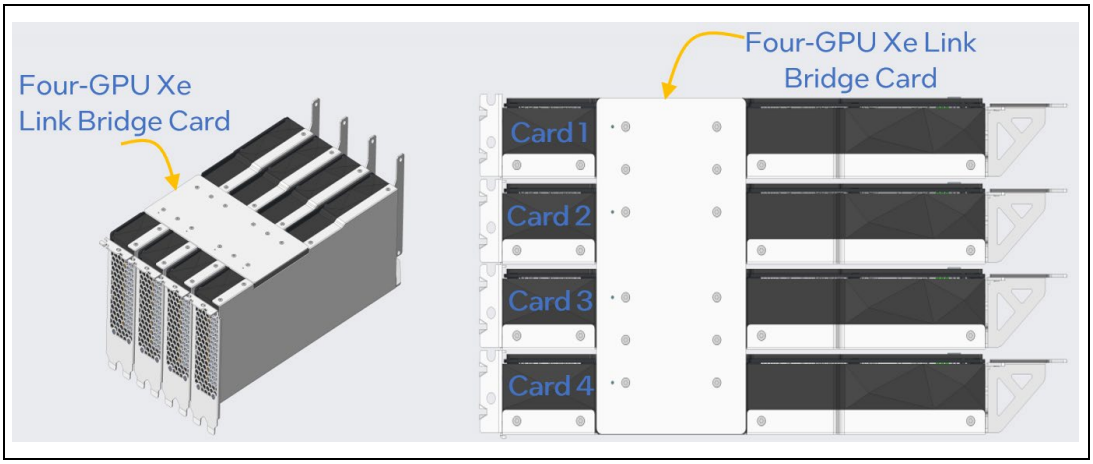

Figure 3-7. Four-Card Intel Xᵉ Link Bridge Connector

4 Intel Data Center GPU Max 1100 Intel Xe Link Bridge Card Installation

This section describes the basic steps and requirements to install Intel Data Center GPU Max 1100 AICs with an Intel Xe Link bridge card into a compatible host server. See *Intel® Xe Link Bridge Card Datasheet*, document number 788941.

4.1 Identify and Prepare Compatible PCIe Slots

In a compatible server, depending on the card configuration, identify one, two, or four adjacent PCIe 5.0 x16 electrical double-wide slots. Confirm each slot can provide adequate airflow for AIC cooling. Install and secure the cards and connect the 12V-2x6 H++ auxiliary power connectors.

4.2 Install the Intel Xe Link Bridge Card

Remove the Intel Xe Link bridge edge connector protective covers from the cards by pulling them straight up and away from the cards' top side.

Carefully orientate the Intel Xe Link bridge card to the connectors shown in the following figure. Firmly and evenly press the bridge card into place until it seats flush with the AICs' top sides.

Figure 4-1. Intel Xe Link Bridge Card before Installation

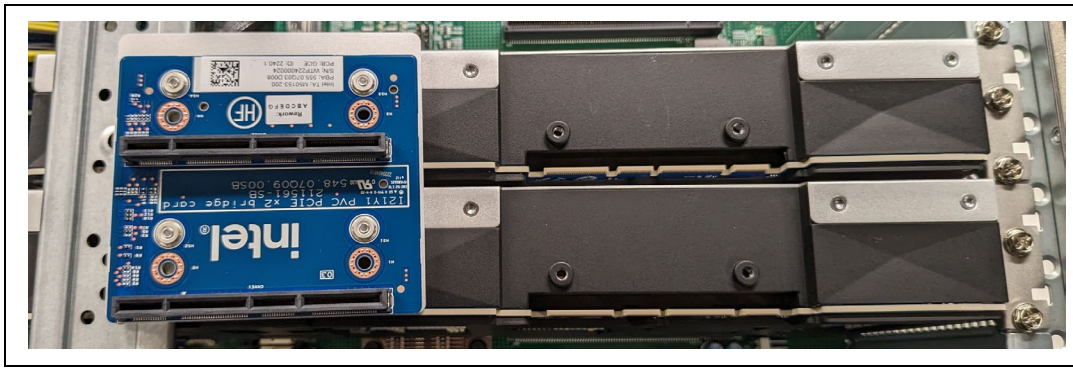

Use a T8 Torx* bit to install the bridge cards' mounting screws in the positions marked with red arrows shown in the following figure. Incrementally tighten the screws using an alternating cross pattern for x2 cards, or for x4 cards as shown in Figure 4-3 below, to the torque specifications indicated in Table 4-1.

Figure 4-2. Intel Xe Link x2 Bridge Card Mounting Screw Locations

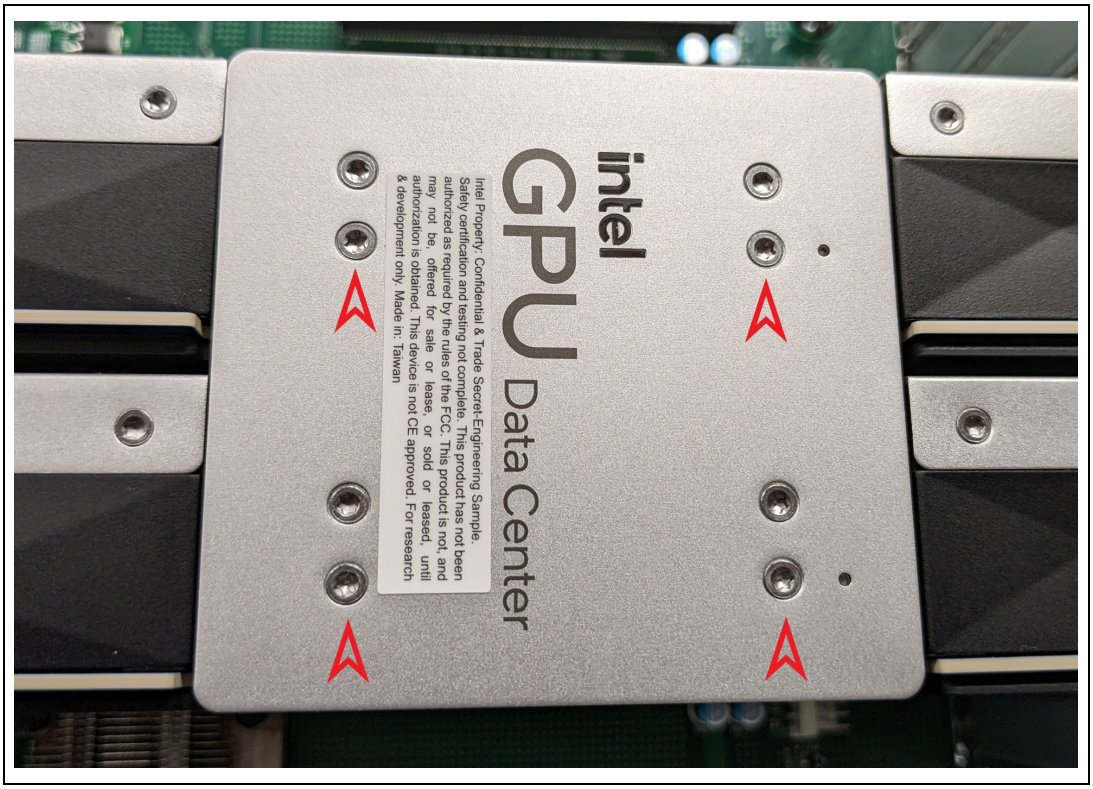

Figure 4-3. Intel Xe Link x4 Bridge Card Screw Tightening Pattern

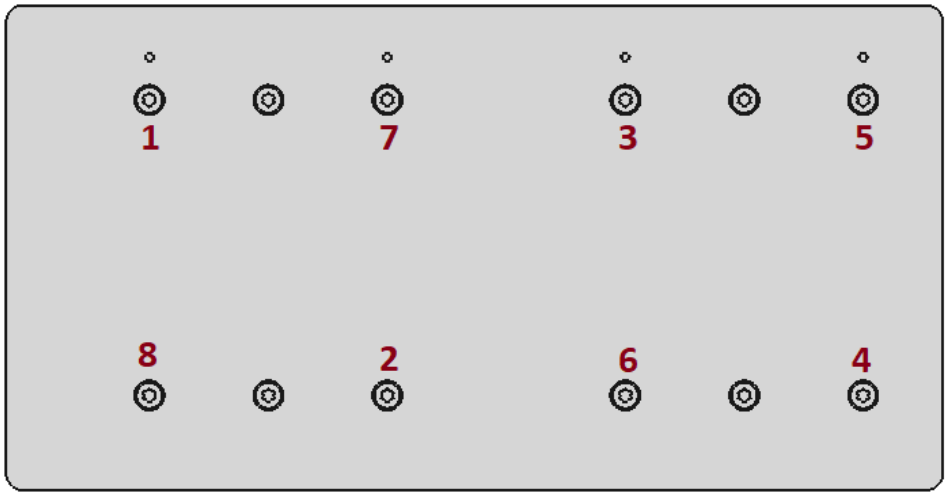

Table 4-1. Intel Xe Link Bridge Card Mounting Screw Torque Specification

Screw	Bit	Torque lbf-in	Tolerance lbf-in	Torque kgf-cm	Tolerance kgf-cm
M3	T8	5	±1	5.76	±1.15

5 Electrical Specifications

5.1 Power Specifications

The Intel Data Center GPU Max 1100 uses high efficiency (efficiency >85%) switched voltage regulators for all on-card power rails, fed by a single CEM 5.1 compliant 12V-2x6 H++ auxiliary power connector.

See Table 5-1 for power delivery requirements.

Table 5-1. Intel Data Center GPU Max 1100 Card Power Requirements

Idle Power	Thermal Design Power (TDP)	Programmable Peak Power	PWRBRK# Asserted
Approximately 50W	300W	1.2 – 2.0 X TDP, default 1.52	95W

Note: Dynamically set power limits will return to default after SOC or host resets.

5.1.1 Power Input Specifications

The AIC TDP is 300W, delivered by a CEM 5.1 compliant 12V-2x6 H++ power connector. See the following table for a detailed power breakdown.

Table 5-2. 12V-2x6 H++ Connector and Board Power Distribution Breakdown

Power Input Source	Voltage (V)	TDC (A)	Icc_max (A)
+12V_EXT_A	12	13.6	27.2
+12V_EXT_B	12	13.6	27.2
+12V_BUS	12	4.7	9.4
+3.3V_Bus	3.3	1	0
+3.3V_AUX	3.3	1	0

5.1.1 PWRBRN_N Emergency Power Reduction

PWRBRK_N, or emergency power reduction, is an active low input signal to the accelerator on PCIe edge finger B30. The host platform can assert PWRBRK_N during emergency events such as a power supply passing its current (I) threshold. For the quickest response from the GPU, PWRBRK_N directly connects to the SoC's PMAXEVENT_N pin. The AMC monitors this signal so it can indicate if a power break occurred.

5.1.2 Optional +3.3 Vaux and SMBus Usage

If present on the host, the Intel Data Center GPU Max 1100 cards can use the +3V3 aux power rail and SMBus interface on the PCIe edge connector. Together, these can be used to remotely manage the card via OOB even when the host is in S5 (soft-off). The card uses a power Multiplexer (MUX) between the +3V3 aux and +3V3 power rails so it can function correctly in systems without standby power.

Note: Although Intel Data Center GPU Max 1100 cards can function in systems without +3v3 aux power or an SMBus interface, some management capability is lost.

5.2 Auxiliary Power Input Connector

A CEM 5.1 compliant auxiliary 12V-2x6 H++ power connector with tin plated contacts is located on the card's upper right rear facing edge, as previously shown in Figure 3-1.

The following table details the auxiliary power connector's pinout.

Table 5-3. PCIe Auxiliary Power Connector Pinout

Pin	Signal
1 to 6	+12V
7 to 8	Ground
9 to 10	Sense Pins To AMC, has on-card 4.7k pullup to `3V3_AUX`. Must connect to ground for card to power on.
11 to 12	Ground
S1	`CARD_PWR_STABLE` (not connected)
S2	`CARD_CBL_PRES#`, has on-card pull down with 4.7k to ground.
S3	SENSE0 To AMC, has on-card 4.7k pullup to 3V3 (See table note).
S4	SENSE1 To AMC, has on-card 4.7k pullup to 3V3 (See table note).

Note: Per the *PCI Express Card Electromechanical Specification Revision 5.0*, SENSE0 and SENSE1 (S3 and S4) the sideband pins tell the connected device how much power that device is allowed to draw. The connector may be able to support much more power, but can be limited for various reasons. The pins must be configured for at least 300W. If both S3 and S4 are left ungrounded, the AMC will not allow the card to power on. See the *12VHPWR CEM specification* from PCI-SIG* for details.

Figure 5-1. 12v Auxiliary Power Receptacle

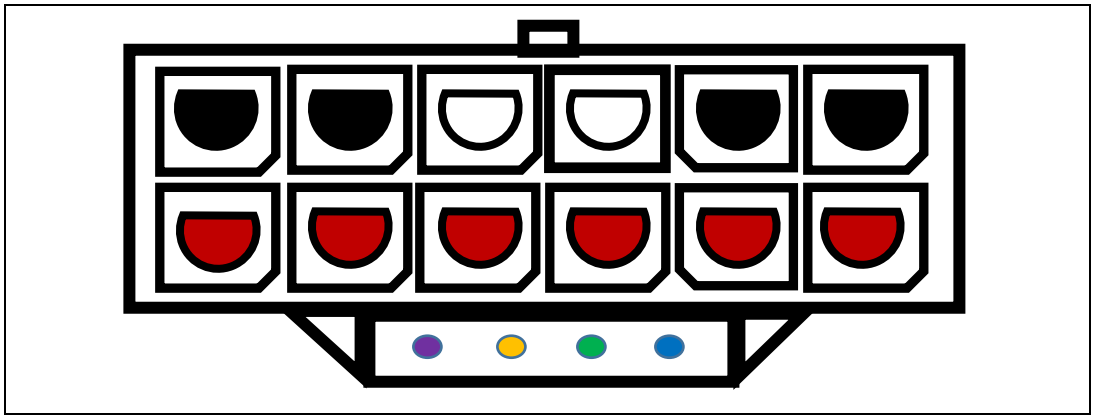

Note: For detailed connector drawings and pinouts, see the *12V-2x6 H++ CEM specification from PCI-SIG*.

5.3 Power Management

5.3.1 Power Capping

The *PCI Express Card Electromechanical Specification Revision 5.0* provides maximum average power rules, but it does not specify peak power and duration. GPU AICs can easily generate peaks exceeding maximum average power specifications. For flexibility, the Intel Data Center GPU Max 1100 offers configurable peak power level control.

A power sensor continually measures instantaneous current, and power drawn from the auxiliary power connector and the 12V PCIe edge connector and is monitored by the CPU over the SVID bus, which throttles the Intel X^{e} stack if set thresholds are exceeded.

A specific maximum TDP, which is maintained by a RAPL mechanism, is fused into the silicon; this can be programmed to a lower value if desired. Short workload bursts are allowed at two higher power levels, known as Power Level 1 (PL1) and Power Level 2 (PL2).

Two power capping loops operate in parallel: a current limiter, as shown in the following figure, and a power limiter. Power capping definitions are described in the following table.

Figure 5-2. Current Limiter in Action

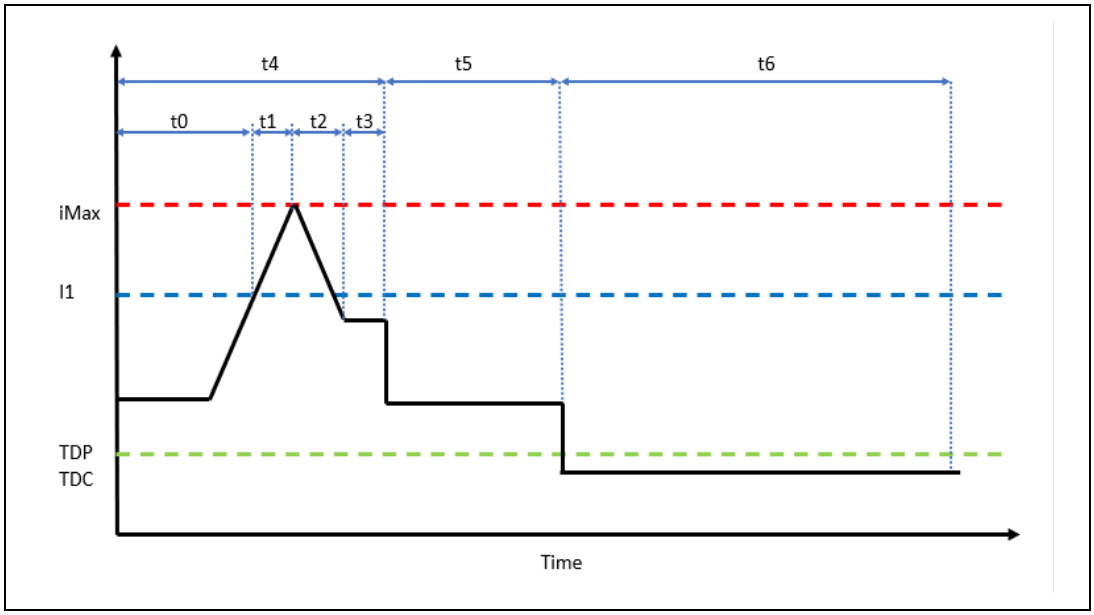

Table 5-4. Power Capping Definitions

Parameter	Description
TDC/TDP	Programmable sustainable current or power, ≤ 300W
I1	Programmable threshold ≤ 2xTDC/TDP. Default = 38A
T0	Time before the breech of I1
t1	About 12 us worst case
t2	About 100 us worst case (default)
t3	Remaining time to t4
t4	500 us
t5	Same as T4 (about 500 us worst case)
t6	t4 + t5 + t6 = TBD (about 8 ms minimum)
Tau1	Time to change the value of the TDP and/or I1. Default = 1 second

Average power will not exceed TDP within a three-Tau window (one second by default; minimum 16 ms). Tau is the user-programmable cycle time during which a spike or pulse can go above TDC/TDP. For example, one pulse above TDC/TDP is allowed every t3 + t4 + t5 ms.

6 Intel Data Center GPU Max 1100 Card Thermal and Mechanical Specifications

6.1 Mechanical Specifications

This section details the Intel Data Center GPU Max 1100 mechanical specifications.

Basic card dimensions are shown in the following table. Dimensioned drawings are shown later in this section.

Table 6-1. Intel Data Center GPU Max 1100 Card Mechanical Dimensions

Parameter	Value
Height	111.15 mm (Full height)
Width	34.35 mm (Dual slot)
Length	266.7 mm
Length with extender bracket	312 mm (Full length)
Card mass (includes I/O bracket)	Approximately 1308g
Card extender bracket mass	35g

6.1.1 Intel Data Center GPU Max 1100 Mechanical Drawings

Figure 6-1. Intel Data Center GPU Max 1100 Card Dimensions

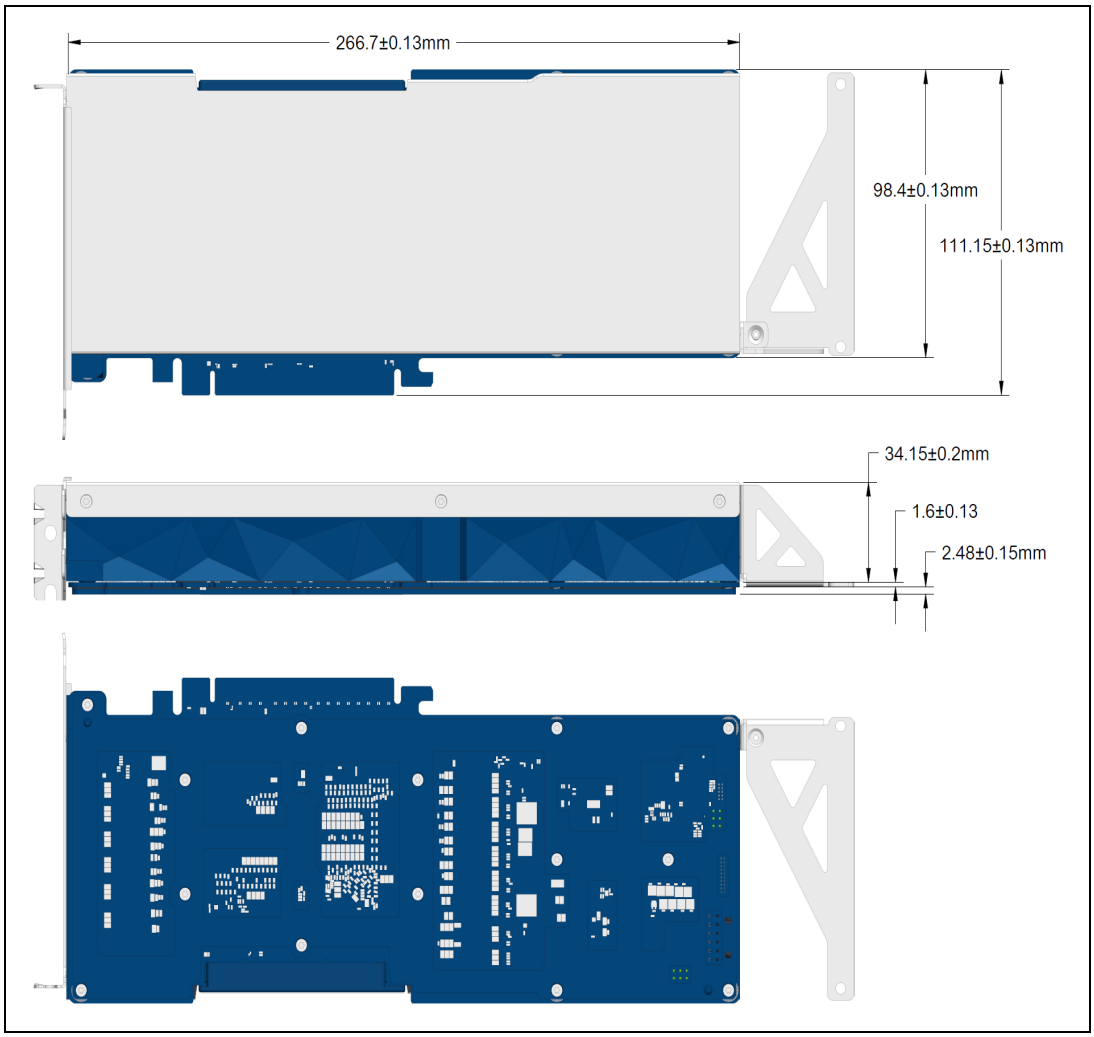

Note: The optional length extension bracket ships detached with the card.

Figure 6-2. Intel Data Center GPU Max 1100 Iso Drawings

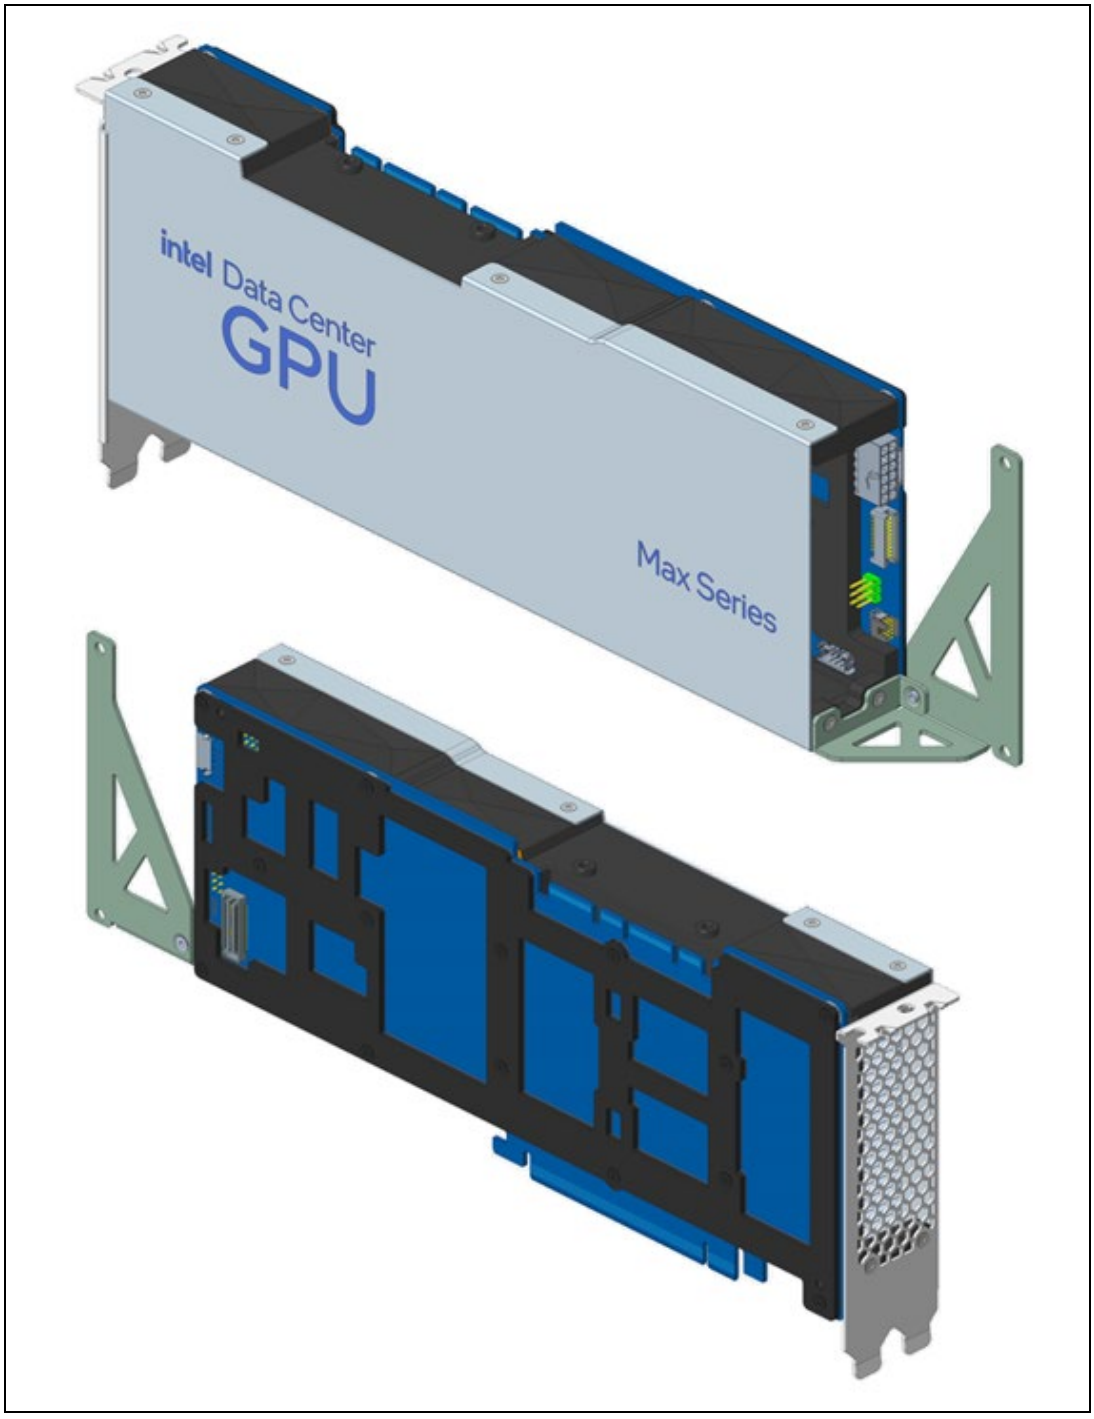

6.2 Thermal Specifications

This section describes the Intel Data Center GPU Max 1100 AIC thermal characteristics.

6.2.1 Airflow Requirements

The graph in the following figure shows the approach air temperature in relation to air speed to the card, as defined for the AIC by the following equation:

$$Y = 0.014 \times x^3 - 0.6694 \times x^2 + 13.324 \times x + 150.78$$

Figure 6-3. Intel Data Center GPU Max 1100 Card Airflow Requirement

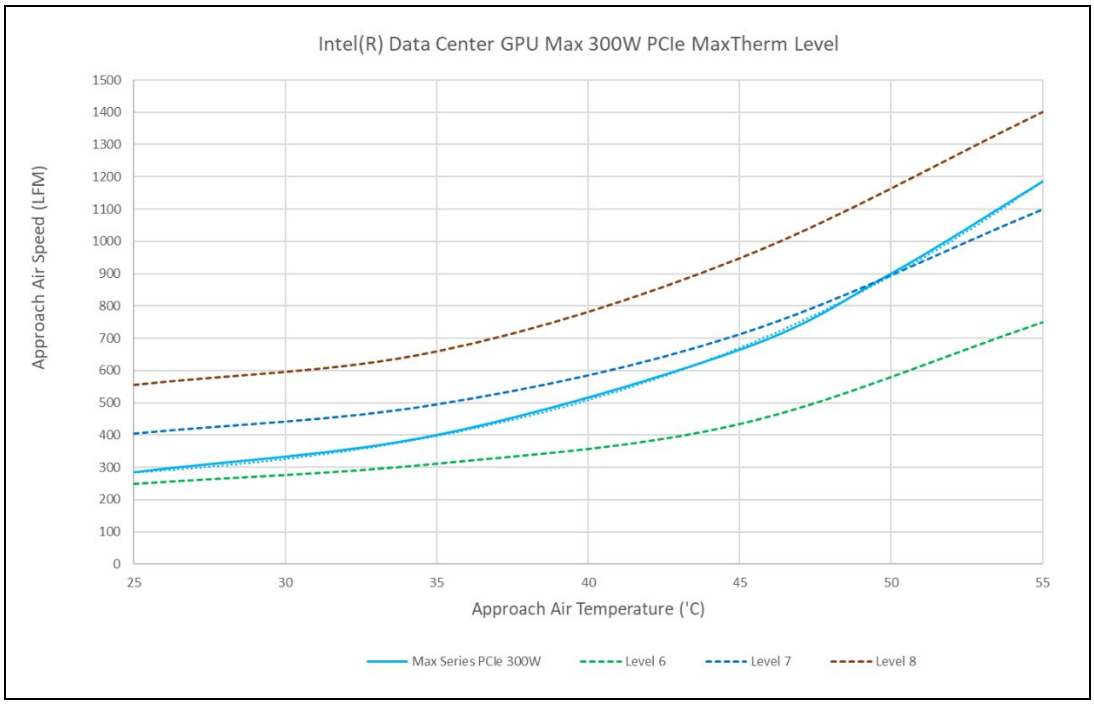

6.2.2 Thermal Management

The GPU PUNIT monitors on-die Digital Temperature Sensors (DTS) to manage the GPU's temperature accordingly by throttling the processing units or HBM when needed.

THERMTRIP_N gets asserted by the GPU PUNIT when either the GPU base die, Intel Xe die, or HBM reaches catastrophic temperatures. When THERMTRIP_N is asserted, the AMC logs the event and shuts down all power rails to the GPU. Once asserted by the PUNIT, THERMTRIP_N remains asserted until one of the following conditions are met:

- VCCINFAON supply is disabled by the AMC.
- PERST_N to the GPU is toggled by the AMC.

VRHOT_N: an active low output from the VR controller, asserted when the VR reaches its temperature threshold. On the AIC, this signal from the VR controller is connected to PROCHOT_IN_N.

PROCHOT_IN_N: active low input to the GPU. When asserted by the VR controller, the die throttles.

The AMC reads the GPU's DTS sensors at regular intervals and reports the values to the host management controller for platform control, such as increasing or decreasing the host system's fan speed. It also reports the card's local ambient temperatures.

6.2.3 Junction Temperature Specifications

Monitoring the die junction temperature is important to understand if the card's thermal boundary conditions are met, as specified in the following table.

Table 6-2. Junction Temperature Specification

Components	Maximum Tjunction (°C)
GPU Thermal Design Target	100
GPU	105
HBM	105

6.2.4 AMC Card Temperature Thresholds

The following table lists the temperature thresholds enforced by the card's AMC.

Table 6-3. Card Temperature Thresholds

Action	Temperature Point (°C)
GPU Warning Threshold	95
GPU Throttle (PROCHOT)	102
Card Shutdown (THERMTRIP)	125
HBM Warning Threshold	95
HBM Reduced Refresh Rate	105

6.2.5 Airflow Impedance Curve

The graph in the following figure shows the airflow impedance curve for the Intel Data Center GPU Max 1100. The curve can be represented by the following equation:

$$y = -0.0137 + 0.0127 * x + 0.0007 * x^2 - 0.0000002 * x^3$$

Figure 6-4. Intel Data Center GPU Max 1100 Airflow Impedance Curve

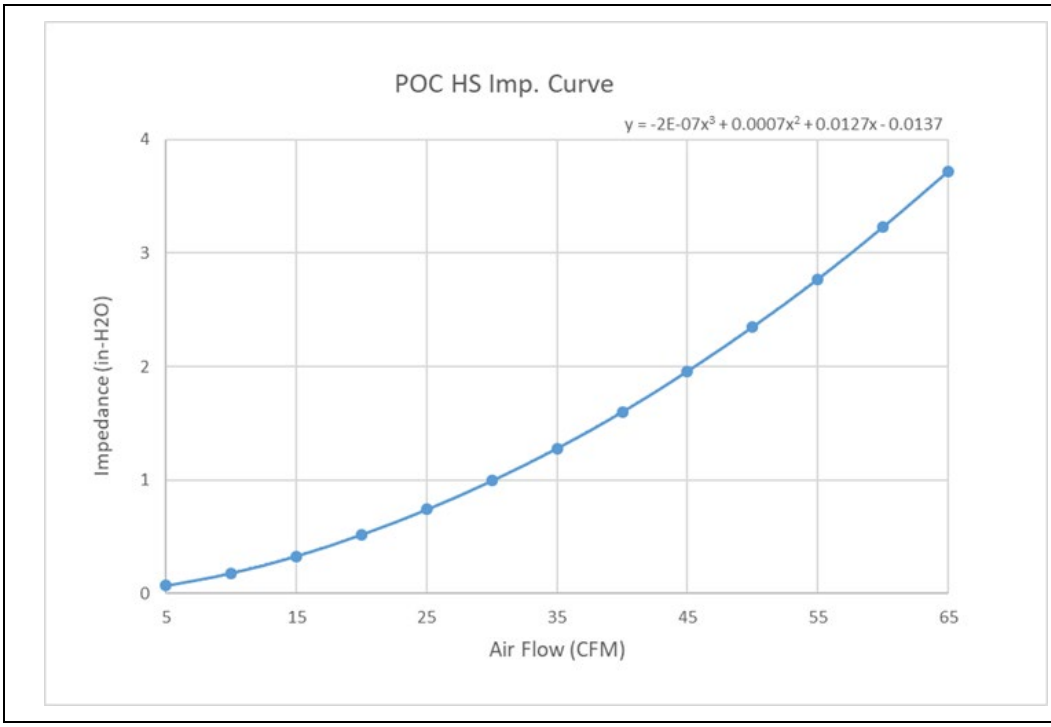

7 Intel Data Center GPU Max 1100 SKUs and Part Numbers

This section contains standard part numbers for the Intel Data Center GPU Max 1100 AIC and the Intel Xe Link bridge cards.

7.1 AIC and Intel Xe Link Bridge Part Numbers

Table 7-1. Card and Intel Xe Link Part Numbers

Part	Intel Part Number	Notes
Intel Data Center GPU Max **1100** Card	99AN5G	AIC and full-length extender bracket
X2 Intel Xe Link Bridge Card	99ARXG	Two-AIC bridge card
X4 Intel Xe Link Bridge Card	99ARXX	Four-AIC bridge card

A Manual AMC Updates and UART Message Monitoring

The Intel Data Center GPU Max 1100 card features a UART for capturing debug information and manual AMC programming for cases where the host server does not have Platform Level Data Model (PLDM), or Intelligent Platform Management Interface (IPMI) enabled for card AMC firmware updates.

Serial update and debug monitoring requires a UART cable, such as a commercially available USB to UART adapter FTDI* TTL-232R-RPI, and AMC updates also require Texas Instruments'* LM Flash Programmer software.

A.1 UART Connections for Manual AMC Programming

To enable AMC program mode, set SW1 position 4 to **0** (**ON**). The switch is located behind the AUX power connector, as shown in the following figure.

Figure 7-1. SW1 DIP Switches (Card Cover Removed for Picture)

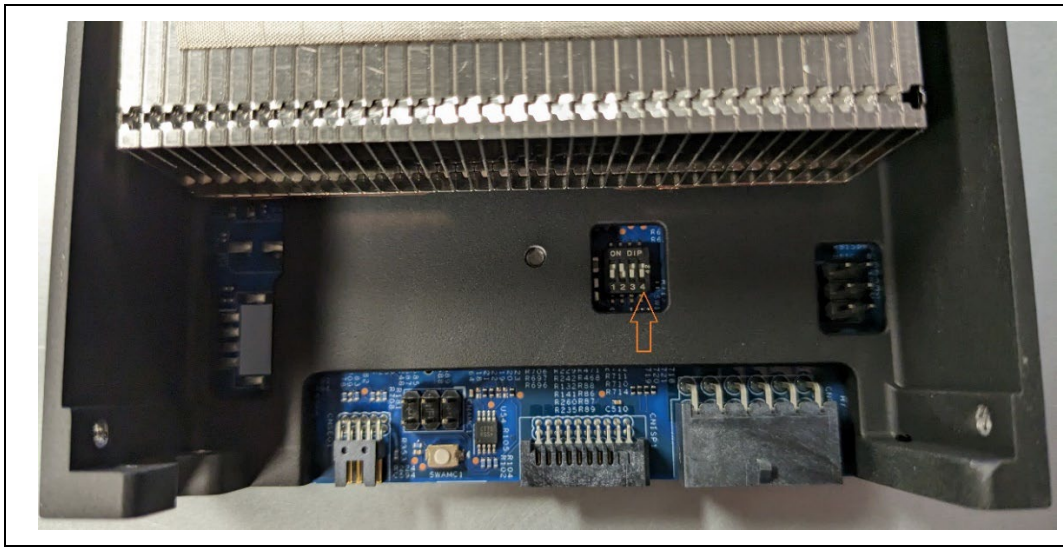

Note: The switch can be accessed without removing the card cover. A pick or similar tool may be required.

Connect a UART cable. The TTL-232R-RPI cable colors are shown as follows:

- RX (yellow, pin 6)
- GND (black, pin 4)
- TX (orange, pin 2)

The following figure shows a connected UART cable.

Figure 7-2. AMC UART Cable Connections

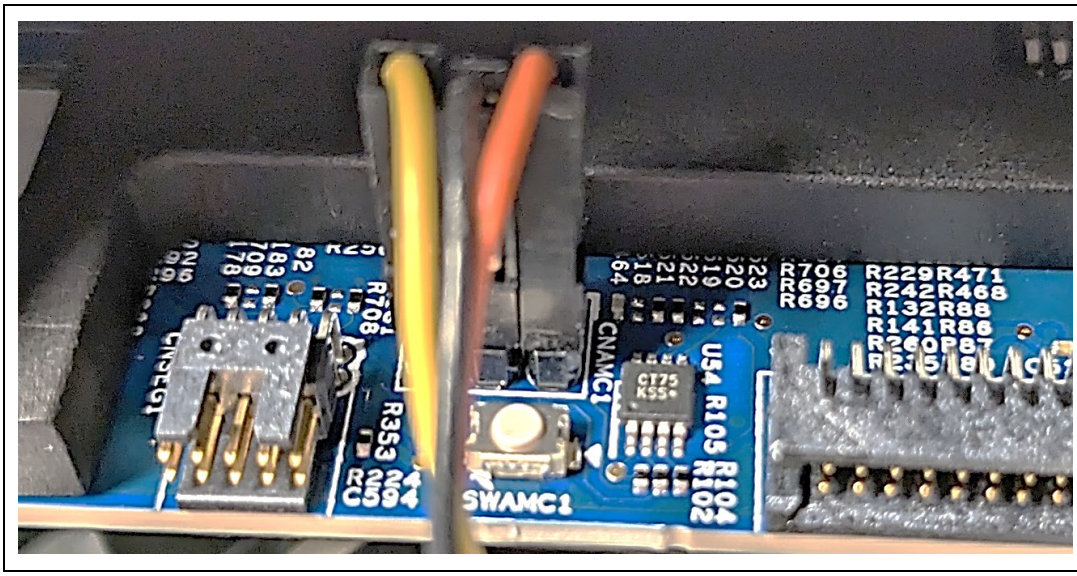

During programming, the card must be powered by the host system.

Use the following settings in the LM Flash Programmer utility:

Configuration Tab:

- Quick Set → Manual Configuration
- Interface → Serial, set COM # to the UART cable port using 115200 Baud.

Program Tab:

- Select an **update** AMC .bin. File → Browse → Select the AMC file to flash.
- Options → Check the box for Reset MCU After Program.
- Options → Program Address Offset: **0x10000**.

Note: The AMC bootloader may be locked. Updating AMC from offset 0x0 with a full image on a locked bootloader will put the AMC into an unworkable state until it is reprogrammed with an appropriate image.

- Select **Program**
- When finished, return SW1 position 4 to **1**.

A.2 AMC Monitoring over UART

To confirm the AMC flash version or to monitor AMC debug output, use a terminal program configured to the UART cable's COM port, 115200, 8n1. No flow control.

Once the message **[36] Waiting for i915 driver load** appears in the com terminal, enter **version** to display the AMC version information.

Figure 7-3. AMC Version Information Shown from Terminal

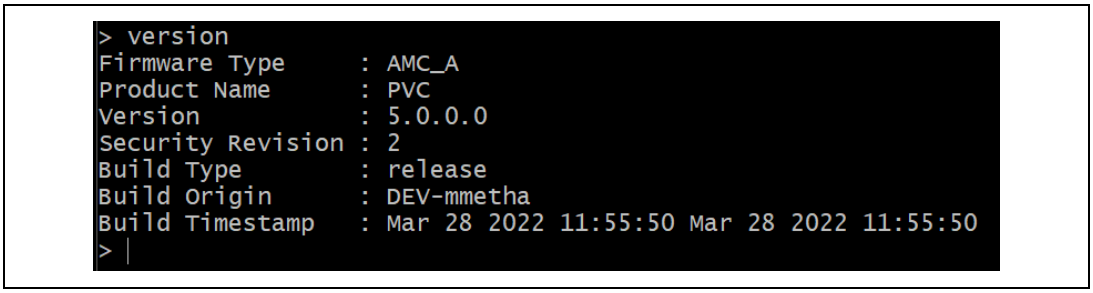

The AMC output can be monitored for debug information. Enter **help** in the terminal to list available commands.

A.3 AMC Status LEDs

The AMC provides card status LEDs, as shown further down in Figure 7-4, which are visible through the I/O bracket vent. For debug, the status LEDs can be used to determine the card's current state during power-up and driver load. If all status BIT LEDs are off and the green AMC heartbeat indicator is blinking, the card booted successfully. See Figure 7-4 for LED positions and locations.

Table A-2. Card Status LEDs

Bit 3	Bit 2	Bit 1	Bit 0	Value	Boot Status
OFF	OFF	OFF	OFF	0	BOOT_SUCCESSFUL
OFF	OFF	OFF	**ON**	1	SOC_BOOT_IN_PPROGRESS
OFF	OFF	**ON**	OFF	2	WAITING_FOR_RESETOUT
OFF	OFF	**ON**	**ON**	3	CSC_REGISTERS_NOT_UPDATED
OFF	**ON**	OFF	OFF	4	PECI_IP_NOT_RESPONDING
OFF	**ON**	OFF	**ON**	5	PCIE_PHI_FW_LOAD_FAILED
OFF	**ON**	**ON**	OFF	6	BOOT_FAIL_UNKNOWN_REASON
OFF	**ON**	**ON**	**ON**	7	HBM_TRAINING_FAILED
ON	OFF	OFF	OFF	8	PUNIT_INIT_FAILED
ON	OFF	OFF	**ON**	9	GT_INIT_FAILED
ON	OFF	**ON**	OFF	10	SOC_INIT_FAILED

Figure 7-4. Debug and Card Status LED Location (I/O Bracket Removed)