NVIDIA Jetson T5000 Modules

Blackwell GPU + Arm® Neoverse® V3AE CPU + LPDDR5X

Data Sheet

DS-11945-001_v1.1 | July 2025

Document History

DS-11945-001_v1.1

Version	Date	Description of Change
1.0	June 18, 2025	Initial release
1.1	July 3, 2025	Updated the following: > Changed 2x DMIC to 1x DMIC

Table of Contents

List of Figures

List of Tables

Chapter 1. Overview

Module	Description
Jetson T5000	Blackwell GPU + Arm Neoverse V3AE CPU + 128GB LPDDR5X

Description		Operation/Performance
AI Performance		
Number of Tensor Core Operations (up to)	Sparse	MAXN: Up to 2070 FP4 TFLOPs |1035 FP8 TFLOPs
		120W: Up to 1820 FP4 TLOPs | 910 FP8 TFLOPs
	Dense	MAXN: Up to 1035 FP4 TFLOPs | 517 FP8 TFLOPs
		120W: Up to 910 FP4 TFLOPs | 455 FP8 TFLOPs
Blackwell GPU		
3GPC, 10 TPC, 2560 NVIDIA® CUDA® cores | 96 5th GEN Tensor cores | MIG Support End-to-end lossless compression | Tile Caching | OpenGL® 4.6+ | OpenGL ES 3.2 | Vulkan™ 1.1+ | CUDA 11.4+		
Operating Frequency (up to):		MAXN: 1.575 GHz
		120W: 1.386 GHz
CUDA Cores (up to):		MAXN: 8.064 FP32 TFLOPs
		120W: 7.096 FP32 TFLOPs
Programmable Vision Accelerator (PVA)		
1x NVIDIA PVA 3.0		
Operating Frequency (up to):		1.215 GHz
PVA Performance (up to):		165 FP32 GFLOPs | 320 FP16 GFLOPs

Description	Operation/Performance
PVA Throughput (up to):	VPUs: 622 INT16 GMAC/s | 2488 INT8 GMAC/s
	PPEs: 194 INT16 GMAC/s | 97 INT32 GMAC/s
CPU Complex	
14 CPU clusters | One core per cluster | Arm® Neoverse V3AE (64-bit) symmetric multi-processing (SMP) CPU architecture | L1 Cache (I, D) per core: 64KB + 64KB | L2 Cache per core: 1MB | L3 Cache 16MB shared System Cache | SPECint@2017_int_base (one core): 6.6 | SPECrate@2017_int_base (all cores): 80	
Operating Frequency per Core (up to):	2.6 GHz
Memory Subsystem	
128GB (8x 16GB) memory with 256-bit LPDDR5X DRAM | 128-bit AES Encryption | System MMU | TrustZone (TZ) Secure and OS-protection regions	
Maximum Operating Frequency (up to)	4,266 MHz
Peak Memory Bandwidth	273 GB/s
HD Video	
Decode 2x NVDEC @ 120W: 1.56 GHz | MAXN: 1.69 GHz NVDEC is on the same rail as the GPU. Supported Standards: H.265 (HEVC), H.264, VP9, VP8, AV1, MPEG-4, MPEG-2, VC-1 (see Multi-Standard Video Decoder section for detailed description)	
Encode 2x NVENC @ 120W: 1.56 GHz | MAXN: 1.69 GHz NVENC is on the same rail as the GPU. Supported Standards: H.265 (HEVC), H.264 (see Multi-Standard Video Encoder section for detailed description)	
Display Controller Subsystem	
4x shared HDMI 2.1 | VESA DisplayPort 1.4a HBR3, MST	
Maximum Resolution DP/HDMI (up to)	7680x4320 (8K) at 30 Hz
PCLK (up to)	1080 MHz
Always On Dual Audio & Sensor Processor Engine	
The Always On (AO) Dual Audio & Sensor Processor Engine (APE/SPE) is a Xtensa F1 DSP(Note 4) subsystem with integrated instruction cache (I-cache), data cache (D-cache) and a tightly coupled memory (TCM) interface	
Audio Subsystem	
Dedicated programmable audio processor | Dual Cadence Tensilica HiFi 5 DSP(Note 4) | PDM in/out | Industry-standard High-Definition Audio (HDA) controller provides a multi-channel audio path to the HDMI® interface	
Networking	
4x up to 25 Gbps MGBE | Media Access Controller (MAC)	

Description	Operation/Performance
Imaging	
Up to six cameras through 16x lanes MIPI CSI-2 | Up to 32 cameras using Virtual Channels | C-PHY 2.1 (10.25 Gbps) | D-PHY 2.1 (40 Gbps)	
Storage	
Internal Storage > 64 MB NOR FLASH Supports External Storage > NVMe through PCIe • Maximum use case of four lanes > SSD through USB 3.2: • USB 3.2 Port 0, 1, or 2	
Peripheral Interfaces	
USB: xHCI host controller with integrated PHY (up to) 3x USB 3.2, 4x USB 2.0	
PCIe: Up to GEN5 (supports up to x8 lanes) | Root Port only – C1 (x1) and C3 (x2) | Root Port or Endpoint - C2 (x1), C4 (x8), and C5 (x4)	
Audio: 5x I2S/2x Audio Hub (AHUB) | Supports I2S, RJM, LJM, PCM, TDM (multi-slot mode) | 1x DMIC	
UART: 4x UART **CAN:** 4x CAN **SPI:** 3x SPI **I2C:** 12x I2C **PWM:** 6x PWM outputs	
Mechanical	
Module Size: 87.0 mm (width) x 100 mm (length) x 15.29 mm (height) | 699 pin B2B Connector | Integrated Thermal Transfer Plate (TTP) with Heatpipe	
Operating Requirements	
Jetson T5000 Power Modes: 75W | 95W | 120W(Note 2) | MAXN(Note 3) Maximum Total Module Power (TMP): 130W	
Temperature Range (at Thermal Transfer Plate (TTP) surface)(Note 1) (T_{TTPMAX}): 80°C | Operating Temperature Range (T_j) (Note 1): -25°C ~ 115°C | Slowdown Temperature (T_j): 109°C | Supported Input Voltage Range (SYS_VIN_HV): 7V-20V | Operating Lifetime (24x7) (Note 1): 5 years	

Notes:

Refer to the Software Features section of the latest *NVIDIA Jetson Linux Development Guide* for a list of supported features; all features may not be available.

1. **See the *Jetson Thor Series Modules Thermal Design Guide* for details.**
2. **Custom power modes can be generated using Jetson Power Estimator. 120W is default power mode.**
3. **GPU/CPU clocks are throttled when TMP exceeds 130W.**
4. **Software not supported.**
5. **FP8 TFLOPs is equivalent to INT8 TOPs.**

Chapter 2. Functional Description

NVIDIA® Jetson® Thor Modules blends industrial-leading performance, power efficiency, integrated deep learning capabilities, and reach I/O to enable emerging technologies with compute-intensive requirements. The Jetson T5000 is designed for a wide variety of applications requiring varying performance metrics.

Figure 2-1. System Block Diagram

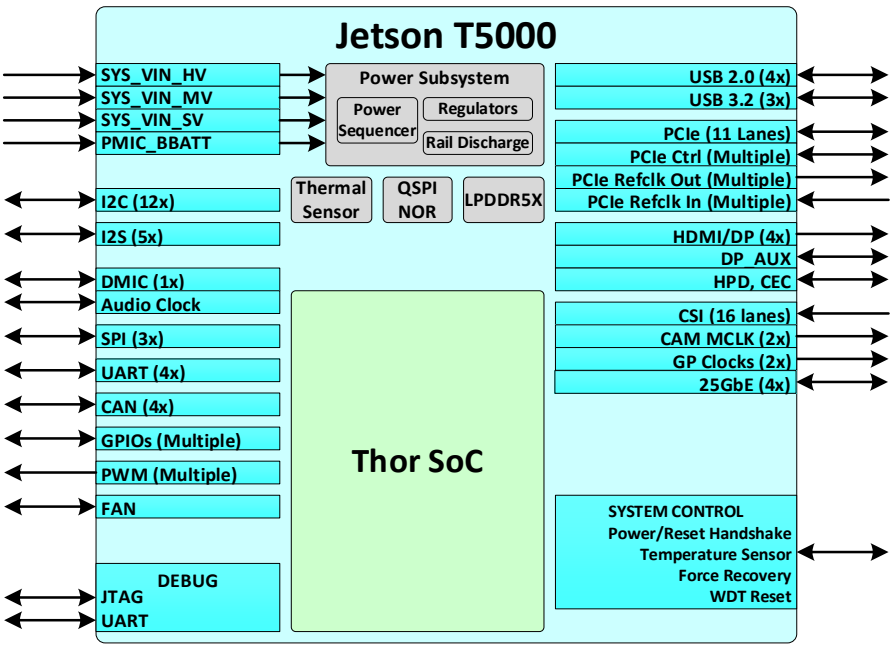

2.1 Blackwell GPU

Highlights of the GPU architecture in Thor SoC:

- Advanced graphics capabilities, with next generation DLSS.
- Fourth generation Tensor Core, adding support for high throughput FP8 compute to reduce the need for FP16 compute.
- Transformer Engine accelerates Transformer model inference.
- Multi-Instance GPU (MIG) support provides a multi-domain computing capability, allowing the GPU to be separated into separate GPU instances where required.

2.1.1 Multi-Instance GPU (MIG)

MIG allows the GPU to be split into two physically separated GPUs to provide isolation, for example allowing allocation of one MIG instance for critical workloads and another MIG instance for other applications. MIG separation is at the GPC level in the GPU.

Graphics can only be supported on one MIG partition, or across the whole GPU, when MIG is not being used. Compute functionality is available on all MIG partitions. Below is an example of MIG partitions:

- MIG_0: 4 TPC (GPC_0), MIG_1: 4 TPC (GPC_1) + 2 TPC (GPC_2)

2.1.2 Transformer Engine

The transformer engine uses a combination of software and custom NVIDIA Tensor Core technology designed specifically to accelerate Transformer model inference.

2.1.3 FP8 and FP4 Support

2.1.3.1 FP8

The FP8 format is normally based on an E4M3 encoding that includes:

- One sign bit, four exponent bits, three mantissa bits
- Subnormals and NaNs are supported

The advantage of the FP8 format is a much wider dynamic range than INT8, so it is less dependent on post training quantization parameters. The dynamic range gain is approximately 2000x.

An E5M2 FP8 mode is also supported. This format is expected to be used for gradients during training, not for inference.

2.1.3.2 FP4

The FP4 format is based on an E2M1 encoding that includes:

- One sign bit, two exponent bits, one mantissa bit
- Infinity and NaNs are not supported

Though the data elements are four bits, they can be scaled by a UE8M0 (for every 32 elements) or UE4M3 (for every 16 elements) scale. Hence, the FP4 format supports a wide dynamic range.

Note: Refer to the Cuda Parallel Thread Execution ISA for additional details.

2.1.4 Compute Features

Blackwell introduces fourth-generation NVIDIA Tensor Cores which offer a wider range of precisions including TensorFloat-32 (TF32), bfloat16, FP16, FP8, FP4, and INT8, all of which provide unmatched versatility and performance.

TensorFloat-32 (TF32) is a new format that uses the same 10-bit Mantissa as half-precision (FP16) math and is shown to have more than sufficient margin for the precision requirements of AI workloads. In addition, since the TF32 adopts the same 8-bit exponent as FP32 it can support the same numeric range.

Blackwell adds support for structured sparsity. Not all the parameters of modern AI networks are needed for accurate predictions and inference; some can be converted to zeros to make the models “sparse” without compromising accuracy. The Tensor Cores in Blackwell can provide up to 2x higher performance for inference of sparse models.

Blackwell supports Compute Data Compression which can accelerate unstructured sparsity and other compressible data patterns. Compression in L2 provides up to a 4x improvement in DRAM read/write bandwidth, up to 4x improvement in L2 read bandwidth, and up to a 2x improvement in L2 capacity.

Blackwell also supports many other enhancements for higher compute throughput.

Table 2-1. GPU Operation

Module	CUDA Cores	Tensor Cores	Operating Frequency per Core (up to)
Jetson T5000	2560	96	1.575 GHz (MAXN)
			1.386 GHz (120W)

2.1.5 Graphic Features

Blackwell graphics capabilities include:

- End-to-end lossless compression, including Post-L2 compression, enabling compression of M stores.
- Tiled Caching
- OpenGL 4.6+, Vulkan 1.2+, CUDA 11.4+
- Adaptive Scalable Texture Compression (ASTC) LDR profile supported
- Modern Graphics features:
  - Ray Tracing
  - DL Inferencing
  - Mesh Shaders
  - Sampler Feedback
  - Variable Rate Shading
  - Texture LOD in compute programs

- Iterated blend, ROP OpenGL-ES blend modes
- 2D BLIT from 3D class avoids channel switch
- 2D color compression
- Constant color render SM bypass
- 2x, 4x, 8x MSAA with color and Z compression
- Non-power-of-2 and 3D textures, FP16 texture filtering
- FP16 shader support
- Geometry and Vertex attribute instancing
- Parallel pixel processing
- Early-z reject: Fast rejection of occluded pixels acts as multiplier on pixel shader and texture performance while saving power and bandwidth
- Video protection region

2.2 Multi-Standard Video Decoder

The module incorporates two instances of the NVIDIA Multi-Standard Video Decoder (NVDEC). This video decoder accelerates video decode, supporting low resolution mobile content, Standard Definition (SD), High Definition (HD) and UltraHD (8K, 4K, etc.) video profiles. The video decoder communicates with the memory controller through the video DMA which supports a variety of memory format output options. For low-power operations, the video decoder can operate at the lowest possible frequency while maintaining real-time decoding using dynamic frequency scaling techniques.

Video decode standards supported: H.265 (HEVC), H.264, VP9, VP8, AV1, MPEG-4, MPEG-2, and VC-1.

Table 2-2. Supported Video Decode Stream Specification

Standard	Profiles	1080p30	1080p60	4Kp30	4Kp60	8Kp30	Mpix/Sec (per instance)	Total Mpix/sec (2x)
120W Mode:								
H264	Baseline, Main, High,	80x	40x	20x	10x	4x	2520	5040
	High 444, High 444 Predictive, MVC (per view considering 2 views)*	40x	20x	10x	4x	2x	1260	2520
	YUV 422	60x	30x	14x	6x	2x	1930	3860
HEVC	Main, Main10	78x	38x	18x	8x	4x	2460	4920

Standard	Profiles	1080p30	1080p60	4Kp30	4Kp60	8Kp30	Mpix/Sec (per instance)	Total Mpix/sec (2x)
	Main444, Main444 10, MV*	38x	18x	8x	4x	2x	1230	2460
	YUV 422	58x	28x	14x	6x	2x	1880	3760
VP9	Profile 0, Profile 2	58x	28x	14x	6x	2x	1840	3780
AV1	Main Profile	48x	24x	12x	6x	2x	1560	3120
MAXN Mode:								
H264	Baseline, Main, High,	84x	42x	22x	10x	4x	2740	5480
	High 444, High 444 Predictive, MVC (per view considering 2 views)*	42x	20x	10x	4x	2x	1370	2740
	YUV 422	66x	32x	16x	8x	2x	2100	4200
HEVC	Main, Main10	82x	40x	20x	10x	4x	2680	5360
	Main444, Main444 10, MV*	40x	20x	10x	4x	2x	1340	2680
	YUV 422	64x	32x	16x	8x	2x	2050	4100
VP9	Profile 0, Profile 2	64x	32x	16x	8x	2x	2000	4000
AV1	Main Profile	54x	26x	12x	6x	2x	1700	3400

Notes:

- **For 444-related profiles, perf numbers given are for YUV444 input; for other profiles it is for YUV420 8-bit input.**
- **The perf data is indicative, based on varied content/bitrate evaluations, and should serve as general performance guidance.**

2.3 Multi-Standard Video Encoder

The module incorporates two instances of the NVIDIA Multi-Standard Video Encoder (NVENC). This multi-standard video encoder enables full hardware acceleration of various encoding standards. It performs high quality video encoding operations for mobile applications such as video recording and video conferencing. The encode processor is designed to be extremely power efficient without sacrificing performance.

Video encode standards supported: H.265 (HEVC), H.264.

Table 2-3. Supported Video Encode Stream Specification

Standard	Profiles	1080p30	1080p60	4Kp30	4Kp60	Mpix/Sec (per instance)	Total Mpix/sec (2x)
120W Mode:							
H264	HQ	14x	6x	2x	0x	445	890
	HP	32x	16x	8x	4x	1040	2080
	UHP	44x	22x	10x	4x	1385	2770
HEVC	HQ	18x	8x	4x	2x	590	1180
	HP	26x	12x	6x	2x	830	1660
	UHP	50x	24x	12x	6x	1585	3170
MAXN Mode:							
H264	HQ	14x	6x	2x	0x	485	970
	HP	36x	18x	8x	4x	1125	2250
	UHP	48x	24x	12x	6x	1500	3000
HEVC	HQ	20x	10x	4x	2x	640	1280
	HP	28x	14x	6x	2x	900	1800
	UHP	54x	26x	12x	6x	1720	3440

Notes:

- The perf numbers given are for YUV420 8-bit input.
- The perf data is indicative, based on varied content/bitrate evaluations, and should serve as general performance guidance.

2.4 Programmable Vision Accelerator

The Thor PVA is the third generation of NVIDIA's vision DSP architecture, which is an application-specific instruction vector processor that targets computer-vision along with virtual and mixed reality applications. These are some key areas where PVA capabilities are a good match for algorithmic domains that need to have a predictable processing capability, at low power and low latency. In addition, the PVA is a good computational engine for sensor processing across various sensors (camera, radar, lidar, etc.) freeing up the GPU for DL intensive tasks, as well as extending battery life and operating temperature range of end products in robotics.

- PVA throughput:
  - VPUs provide up to 622 INT16 GMAC/s, 2488 INT8 GMAC/s
  - Up to 165 FP32 GFLOPs | 320 FP16 GFLOPs
  - PPEs provide up to 194 INT16 GMAC/s, 97 INT32 GMAC/s
- A PVA cluster has the following components:
  - Dual Vector Processing Units (VPU) with vector cores, instruction cache, and vector data memories. The VPU has a single instruction stream with Very Long Instruction Word (VLIW) slots for both scalar and vector instructions.
  - Dual 2D Pixel Processing Engine (PPE) to VPU for additional pixel processing capabilities with additional MACS available.
  - Dual Decoupled Lookup Unit (DLUT) for concurrent lookup and interpolation of VPU.
  - Dual DMA engines with multi-dimensional tensor access through abstracted data flow API access provided by PVA SDK.
  - 512 KB of local memory (VMEM) per VPU divided into 4x 128 KB super-banks
  - 1 MB of L2SRAM
  - Dedicated hardware offload for task scheduling and control.

2.5 Optical Flow Accelerator

The Optical Flow Accelerator (OFA) is a hardware accelerator for computing optical flow and stereo disparity between the frames.

OFA can operate in Stereo Disparity Mode and Optical Flow Mode.

OFA generates disparity and flow vector block-wise, one output for each input block of 4x4, 2x2, and 1x1 pixels (referred as output grid size). The generated output can be further post-processed to improve accuracy, up sampled to produce dense map.

- **Stereo Disparity Mode**
  - OFA processes rectified left and right view of stereo captures and generates disparity values between them.

- The output stereo disparity format is fixed signed 10.5 (2 bytes per disparity output). We need to divide the output values by 32 to get a disparity value in terms of pixel units.

> **Optical Flow Mode**

- OFA generates optical flow between two given frames.
- The input to OFA in this mode is image pyramid of input and reference frames with fixed scale factor of 2. As search range of single layer is small, each pyramid level will search around output of previous pyramid level.
- OFA generates a flow vector has X and Y component that represent motion in X and Y direction. The output flow format is fixed signed 10.5 (4 bytes per flow vector). We need to divide the output values by 32 to get a disparity value in terms of pixel units.

Table 2-4. Optical Flow Accelerator

OFA Parameter	Description
Input Image Size	Minimum size: 32 × 32 Maximum size: 8192 × 8192 No alignment requirement
Input Image format / bit depth	Luma / Single channel Input Supports bit depth of 8/10/12/16 bits
Disparity Output	Disparity Map in fixed S 10.5 format
Flow Output	Flow Map (mvx, mvy) in fixed S10.5 format
Hardware Cost Output	Hardware cost for winner disparity candidate 8 bit per output / Range 0 - 255
Output Grid Size	1x1/2x2/4x4
Maximum Disparity Range	128 / 256
Search Direction	Left / Right Disparity Map
Region Of Interest Support	Supports maximum 32 ROI per stereo pair
Maximum Pyramid Levels for Flow	7

Table 2-5. Optical Flow Accelerator Streams

Mode	Grid Size	Resolution (Maximum Number of Streams)	Throughput (Up to)
120W Mode:			
Optical Flow	4x4	4K30 (4) | 1080p60 (9) | 1080p30 (18)	1100 MP/s
Stereo	4x4	4K30 (6) | 1080p60 (12) | 1080p30 (24)	1455 MP/s
MAXN Mode:			
Optical Flow	4x4	4K30 (4) | 1080p60 (9) | 1080p30 (19)	1180 MP/s
Stereo	4x4	4K30 (6) | 1080p60 (12) | 1080p30 (25)	1550 MP/s

2.6 NVJPEG

The JPEG processing block is responsible for JPEG (de)compression calculations (based on JPEG still image standard), image scaling, decoding (YUV420, YUV422H/V, YUV444, YUV400) and color space conversion (RGB to YUV).

It consists of hardware engine with two instances of NVJPEG HW:

- 2x NVJPEG
- Performance: 2x 1500 MPix/Sec

Table 2-6. NVJPEG Streams per Instance

NVJPEG	Compression Ratio	Format	Throughput (Up to)	Number of 1080p30 Streams	Number of 4Kp30 Streams
120W Mode:					
JPEG Decode	6:1	YUV420	1507 MPix/s	24x	6x
	10:1	YUV420	1775 MPix/s	28x	7x
JPEG Encode	6:1	YUV420	1818 MPix/s	29x	7x
	10:1	YUV420	2348 MPix/s	37x	9x
MAXN Mode:					
JPEG Decode	6:1	YUV420	1617 MPix/s	25x	6x
	10:1	YUV420	1905 MPix/s	30x	7x
JPEG Encode	6:1	YUV420	1951 MPix/s	31x	7x
	10:1	YUV420	2520 MPix/s	40x	10x

Note: 2x NVJPEG engines are present in Thor. The above data is for single instance of NVJPEG.

2.7 CPU Complex

The CPU cluster is comprised of 14 Arm Neoverse V3AE cores. Neoverse features:

- Arm V9.2-A 64-bit CPU architecture
- Dual 128-bit NEON engines
- Dual 128-bit SVE2 engines

Table 2-7. CPU Operation

Module	CPU Cores	CPU Maximum Frequency
Jetson T5000	14	2.6 GHz

Features:

- Armv9.2-A A64 Instruction set
- Scalable Vector Extension (SVE) and Scalable Vector Extension 2 (SVE2)
- 48-bit Physical Address and 48-bit Virtual Address
- Cryptographic Engine for crypto function support
- L1 caches – separate 64 KB I-cache and 64 KB D-cache per core
- L2 cache – a unified, 8-way set associative, 1 MB unified L2 cache per core
- Integrated execution unit with Advanced Single Instruction Multiple Data (SIMD) and floating-point support
- Performance Monitoring Unit (PMU)
- Embedded Trace Extension (ETE)
- TRace Buffer Extension (TRBE)
- Statistical Profiling Extension (SPE)
- Error detection and reporting
- Memory Management Unit (MMU)
- CoreSight (v3.0) support
- Dynamic C7 support
- Dynamic CC7 Power Gating
- Generic Interrupt Controller (GIC) CPU interface to connect to an external interrupt Distributor
- Generic Timers interface supporting 64-bit count input from an external system counter

Performance Monitoring:

- The Hardware Performance Monitor (HWPM) is an NVIDIA function implemented outside of the CPU core. The HWPM uses counters and associated logic to gather various statistics on the operation of the processor and memory system during runtime.

Table 2-8. CPU Cache Configuration

Cache	Capability
L1	64 KB Instruction Cache per CPU core 64 KB Data Cache per CPU core
L2	1 MB Unified Cache per CPU core
L3	16 MB Shared System Cache

2.8 Memory Subsystem

128GB 256-bit LPDDR5X DRAM is used on the Jetson T5000. It supports the following:

- Secure external memory access using TrustZone technology
- System MMU
- Maximum operating frequency: 4,266 MHz
- Low latency accesses for the CPU, whose performance is very latency sensitive
- Guaranteed bandwidth for real time streaming clients like CSI or Ethernet Audio Video Bridging (AVB)
- High sustained bandwidth for clients that are bandwidth sensitive, but latency tolerant such as the GPU
- Virtual channels to avoid blocking across traffic classes

The Memory Subsystem (MSS) provides access to local DRAM, SysRAM, and provides a SyncPoint Interface for inter-processor signaling. The MSS supports full-speed I/O coherence by routing requests through a scalable coherence fabric. It also supports a comprehensive set of safety and security mechanisms.

Structurally, the MSS consists of:

- MSS Data Backbone - routes requests from clients to the MSS Hub and responses from MSS Hub to the clients.
- MSS Hub - receives and arbitrates among client requests, performs SMMU translation, and sends requests to MCF.
- Memory Controller Fabric (MCF) - performs security checks, feeds I/O coherent requests to the Scalable Coherence Fabric (SCF), and directs requests to the multiple memory channels.
- Memory Controller (MC) Channels - row sorter/arbiter and DRAM controllers.
- DRAM I/O - channel-to-pad fabric, DRAM I/O pads, and PLLs.

Jetson T5000 provides three independent column address bits to each sub-partition, allowing it access different 32-byte sectors of a Group of Bytes (GOB) between the sub-partitions. It provides connections between a wide variety of clients, supporting their bandwidth, latency, quality-of-service needs, and any special ordering requirements that are needed. The MSS supports a variety of security and safety features and address translation for clients that use virtual addresses.

Features:

- LPDDR5X: x64 DRAM chips
- 256-bit wide data bus.
  - 32 chip selects
  - Operates in either single or dual rank mode
  - Supports per-byte data masks
- Low latency path and fast read/response path support for the CPU complex cluster.
- Support for low-power modes:

- Software controllable entry/exit from: deep sleep mode and self-refresh
- Hardware dynamic entry/exit from: active power down and self-refresh.
- Disable unused address/command taps based on mode
- Pads use DPD mode during idle periods.

> High-bandwidth interface to the integrated Blackwell GPU.
> Full-speed I/O coherence with bypass for Isochronous (ISO) traffic.
> System Memory-Management Unit (SMMU) for address translation based on the Arm SMMU-500.
> High-bandwidth PCIe ordered writes.
> AES-XTS encryption with 128-bit key.

2.9 Storage

Storage is supported externally via NVMe.

2.10 Video Input

The Video Input (VI) block receives data from the CSI receiver and prepares it for presentation to system memory or the dedicated image signal processor execution resources. The VI block provides formatting for RGB, YCbCr, and raw Bayer data in support of several camera user models. These models include single and multi-camera systems, which may have up to six active streams. The input streams are obtained from MIPI compliant CMOS sensor camera modules.

2.10.1 Image Signal Processor (ISP)

Jetson T5000 has two ISPs. The ISP module takes data from the VI/CSI module or memory in raw Bayer format and processes it to YUV output. The imaging subsystem supports raw (Bayer) image sensors up to 24 million pixels. Advanced image processing is used to convert input to YUV data and remove artifacts introduced by high-megapixel CMOS sensors and optics with up to 30-degree CRA.

The ISP in Jetson T5000 supports a max frequency of 1215 MHz, with a maximum throughput of 3.5 GPixel/s.

Features:

> Flexible post-processing architecture for supporting custom computer vision and computational imaging operations.
> Hardware noise reduction
> Black-level compensation
> Lens-shading compensation
> Bad pixel correction

- Edge enhancement
- Color and gamma correction
- Global and local tone mapping
- Color-space conversion (RGB to YUV)
- Improved de-mosaic algorithms
- Multi Exposure and Piecewise Linear HDR Processing support with maximum 20-bit bit-depth dynamic range

2.10.2 Video Image Compositor (VIC)

VIC implements various 2D image and video operations in a power-efficient manner. It handles various system UI scaling, blending, rotation operations, video post-processing functions needed during video playback, and advanced de-noising functions used for camera capture.

Features:

- HDR Pixel Processing
- Pixel Format Processing
- Temporal Noise Reduction
  - 5x5 Spatial Bilateral Filter de-noising on both luma and chroma channels
  - Programmable domain and ranged filter coefficients
  - Motion-adaptive temporal noise reduction
  - Flexible pixel format support
  - Alpha/Beta rounding and output rounding enhancement
  - Fast alpha recovery to help suppress ghosting side effect
- Scaling
- Color Conversion
- Memory Format Conversion
- Blend/Composite
- 2D Bit BLIT operation
- Rotation
- Geometry Transform Processing
  - Programmable nine-points controlled warp patch for distortion correction
  - Real-time on-the-fly position generation from sparse warp map surface
  - Pincushion/barrel/moustache distortion correction
  - Distortion correction of 180- and 360-degree wide FOV lens
  - Scene perspective orientation adjustment with IPT
  - Full warp map capability
  - Non-fixed Patch size with 4x4 regions

2.11 Dual Audio and Sensor Processing Engine

The Xtensa F1 DSP processor in the Always On (AON) block is also referred to as the Tensilica LX7 with DSP co-processor for Audio and Sensor Processing Engine (APE/SPE). The AON cluster provides all the necessary hardware features to support low power sensor management and wake use cases. The cluster consists of a Cadence Xtensa F1 DSP processor, along with an internal DMA, Control Fabric, Low Speed I/O controllers, and other miscellaneous IPs.

AON Xtensa F1 DSP implementation:

- Armv7-R ISA
- Integrated instruction and data caches
- Tightly coupled memory (TCM) interface for local SRAM
- Vectored interrupt support
- 64-bit AXI Initiator interface for DRAM requests
- 64-bit AXI Initiator interface for MMIO requests
- 32-bit AHB Initiator interface for Arm Vectored Interrupt Controller (AVIC) access
- 64-bit AXI Target interface for DMA access to the local SRAM

Note: Currently not supported in software.

2.12 Display Controller

The NVIDIA Jetson T5000 provides 4x shared HDMI and DP ports. The HDMI and VESA DisplayPort (DP) interfaces share the same set of interface pins.

HDMI provides a unified method of transferring both audio and video data. The HDMI block receives video from either display controller and audio from a separate high-definition audio (HDA) controller; it combines and transmits them as appropriate.

Supported HDMI features are:

- Compliant to HDMI 2.1 (up to 8K30).
  - Support 8/10 bpc RGB or YUV444
- HDCP 2.2 and 1.4
- Multichannel audio from HDA controller, up to eight channels 192 kHz 24-bit.
- 24-bit RGB and 24-bit YUV444 (HDMI) pixel formats.

VESA DisplayPort (DP) is a digital display interface often used to connect a video source to a display device over a cable, in consumer or commercial applications.

Supported DisplayPort features are:

- Compliant to the DisplayPort 1.4a Specification

  - Support 18 bpp RGB
  - Support 24 bpp RGB
  - Support 30 bpp RGB
- Support up to 1080 MHz display clock
- Support for 1/2/4 lanes
- Support for following bit rates:
  - RBR (Reduced Bit Rate, 1.62 Gbps)
  - HBR (High Bit Rate, 2.7 Gbps)
  - HBR2 (High Bit Rate 2, 5.4 Gbps)
  - HBR3 (High Bit Rate 3, 8.1 Gbps)
- Multi-Stream Transport (MST)
- Support for two to eight channels of audio streaming up to 96 kHz sample rate

Notes:

1. HDCP is disabled by default. Contact your NVIDIA support team for activation.
2. Refer to the Software Features section of the latest *NVIDIA Jetson Linux Development Guide* for a list of supported features; all features may not be available.

2.13 High-Definition Audio Subsystem

Standard
High-Definition Audio Specification Version 1.0a

The HD Audio Subsystem uses a collection of functional blocks to off-load audio processing activities from the CPU complex, resulting in fast, fully concurrent, and highly efficient operation. This subsystem is comprised of the following:

- Audio Processing Engine (APE)
- High-Definition Audio (HDA)

2.13.1 Audio Processing Engine (APE)

The Audio Processing Engine (APE) in the SoC can perform audio processing with minimal supervision from the main CPUs if required. It can also split audio processing tasks with the main CPUs, in which case the APE is generally used for more latency critical processes.

The APE contains an Audio DSP cluster (ADSP), an Audio Hub (AHUB) equipped with multiple hardware accelerators for audio signal pre-processing and post-processing, and an Audio DMA engine (ADMA).

Table 2-9. Audio Processing Engine (APE) Features

Features	Capability
Engine(Note)	Dual Cadence Tensilica HiFi 5 DSP cores in separate audio subsystem (APE)
Audio Hub (AHUB) I/O Modules	5x I2S Audio Hub (AHUB) Internal Modules
High-Quality Async Sample Rate Converter (ASRC)	Yes
Low Latency Sample Frequency Converter (SFC)	Yes
Digital Mixer	Digital Audio Mixer: 10-in/5-out: > Up to eight channels per stream > Simultaneous multi-streams > Flexible stream routing
Audio Demuxing and Muxing	Yes
Master Volume Controller	Yes
Parametric Equalizer	Up to 12 bands

Note: Currently not supported in software.

2.13.2 High-Definition Audio (HDA)

Standard
Intel High-Definition Audio Specification Revision 1.0a

The Jetson T5000 implements an industry-standard High-Definition Audio (HDA) controller. This controller provides a multi-channel audio path to the HDMI interface. The HDA block also provides an HDA-compliant serial interface to an HDA codec. Multiple input and output streams are supported.

Features:

- Supports HDMI 2.1 and DP1.4
- Support up to four audio streams for use with HDMI/DP
- Supports striping of audio out across 1,2,4[a] SDO lines
- Supports DVFS with maximum latency up to 208 µs for eight channels
- Audio Format Support
  - Uncompressed Audio (LPCM): 16/20/24 bits at 32/44.1/48/88.2/96/176.4/192[b] kHz
  - Compressed Audio format: AC3, DTS5.1, MPEG1, MPEG2, MP3, DD+, MPEG2/4 AAC, TrueHD, DTS-HD
    - [a] Four SDO lines: cannot support one stream, 48 kHz, 16-bits, two channels; for this case, use a one or two SDO line configuration.

Chapter 3. Power and System Management

3.1 Input Power

The NVIDIA Jetson T5000 has three power inputs:

- SYS_VIN_HV: 7V – 20V
- SYS_VIN_MV: 5V
- SYS_VIN_SV: 3.3V
- PMIC_BBATT: (1.85V to 5.5V input) for RTC backup

Power is then supplied to the devices on board through power sequencer IC and dedicated voltage regulators. All internal module voltages and I/O voltages are generated from these inputs. Input powers must be supplied by the carrier board that the Thor Module is designed to connect to.

PMIC_BBATT provides power for RTC backup.

The input voltage measured at the module connector should never exceed the voltage range defined in Table 3-1.

See the *Jetson Thor Series Modules Design Guide* for details on connecting to each of the interfaces.

Table 3-1. Power and System Control Pin Descriptions

Pin Name	Description	Direction	Pin Type
SYS_VIN_HV	**System Voltage Input High:** input voltage to the module of 7V ~ 20V.	Input	7V to 20V
SYS_VIN_MV	**System Voltage Input Medium**: Input voltage to the module of 5V	Input	5.0V
SYS_VIN_SV	**System Voltage Input Small**: Input voltage to the module of 3.3V	Input	3.3V
PMIC_BBATT	**Real Time Clock**. Optionally used to provide back-up power for the RTC in the Power Sequencer. Connects to a Lithium Cell or similar power source Which provides power to RTC when system is	Input	1.85V to 5.5V

Pin Name	Description	Direction	Pin Type
	disconnected from power. Rechargeable cells or super capacitors are not supported.		
SYS_RESET_N	**System Reset**: Asserted by power sequence in the module during power-on. Can be asserted by the carrier board to reset the SoC, Boot devices, etc., but not the power sequence.	Bidir	Open-drain 1.8V
VDDIN_PWR_BAD_N	**VDD_IN Power Bad**: Carrier board asserts low when SYS_VIN_HV and SYS_VIN_MV are not valid. Carrier board should stop asserting this signal and allow the signal to be pulled high by the internal pull-up on the module when SYS_VIN_HV and SYS_VIN_MV are stable and have reached required voltage levels. This prevents the SoC from powering up until the main input supply voltages are stable. It also can be driven by module if critical condition makes this necessary, such as thermal shutdown. A 10kΩ pull-up to 5V is on the module	Bidir	Open-drain 5.0V
MODULE_POWER_ON	**Module Power On**: Signal asserted to the module to start power-on sequence. This signal should be driven low by the carrier board initially and then driven high when the module is to be powered on.	Input	CMOS – 5.0V
CARRIER_POWER_ON	**Carrier Power On**: Asserted by the module when module power up sequence is complete. This signal low indicates that it is safe for the carrier board to power up. A 10kΩ pull-up to 3.3V is on module.	Output	Open-drain 3.3V
PERIPHERAL_RESET_N	**Peripheral Reset**: Driven from carrier board and AND'd with SYS_RESET_N to drive the Thor SF_SYS_RST_N pin. When PERIPHERAL_RESET_N is asserted, the SoC and QSPI are reset (not Power Sequencer). A 10kΩ pull-up to 1.8V is on the module.	Input	CMOS – 1.8V
FORCE_RECOVERY_N	**Force Recovery strap pin**. Held low when SYS_RESET_N goes inactive (power-on or reset button press) to enter force recovery mode.	Input	CMOS – 1.8V
SLEEP_REQ_N	**Sleep Request**: Requests module to enter Deep Sleep mode (SC7 state). A 10kΩ pull-up to 1.8V is on the module.	Input	CMOS – 1.8V
MODULE_SLEEP_N	**Sleep Acknowledge**: Indicates module is in Deep Sleep mode (SC7 state).	Output	CMOS – 1.8V
THERMAL_ALERT_N	**Thermal Alert**: Thermal alert assertion. The module will throttle to cool down. System needs to actively cool down the module once the signal is asserted low.	Output	Open-drain 3.3V
SYSTEM_OC_N	**System Over-current Thermal warning**: Optionally system asserts the signal to inform module that is at over-current situation.	Input	COMS – 1.8V (note 2)

Pin Name	Description	Direction	Pin Type
MODULE_SHDN_N	**Module Shutdown indicator**: Used to inform carrier board that a shutdown request has occurred on-module (thermal shutdown). A 10kΩ pull-up to 1.8V is on the module.	Output	CMOS – 1.8V (note 2)
PRSNT0	**Present 0**: Tied with PRSNT1 on module. Used to detect when module is connected to the carrier board. Can be used to keep carrier board from powering the module until the module is installed fully in the carrier board. Tied to GND on carrier board if implemented to match reference design.	N/A	N/A
PRSNT1	**Present 1**: Tied to one side of power button on carrier board. Refer to the Jetson T5000 Carrier Board Reference Desing files for the detailed connection.	N/A	N/A
MID0	**Module ID** #0	Not connected	N/A
MID1	**Module ID** #1	Tied to GND	N/A

3.2 Power Rails

SYS_VIN_HV, SYS_VIN_MV, and SYS_VIN_SV must be supplied by the carrier board that the Jetson T5000 is designed to connect to. All Jetson T5000 interfaces are referenced to on-module voltage rails and no I/O voltage is required to be supplied to the module. See the *Jetson Thor Series Modules Design Guide* for details of connecting to each of the interfaces.

3.3 Power Monitoring

Jetson T5000 has a single three-channel INA that can measure power of the module's main power inputs.

3.4 Power Management Controller (PMC)

The PMC power management features enable both high speed operation and very low-power standby states. The PMC primarily controls voltage transitions for the SoC as it transitions to/from different low power modes; it also acts as a target receiving dedicated power/clock request signals as well as wake events from various sources (e.g., SPI, I2C, RTC, USB) which can wake the system from deep sleep state. The PMC enables aggressive power-gating capabilities on idle modules and integrates specific logic to maintain defined states and control power domains during Deep Sleep mode.

3.5 Resets

If SYS_RESET_N is asserted, the Jetson T5000 and onboard storage will be reset. This signal is also used for baseboard power sequencing.

3.6 PMIC_BBATT

An optional back up battery can be attached to the PMIC_BBATT module input to maintain the module real-time clock (RTC) when VIN is not present. Batteries can be used to power the pin, but charging is not supported. This pin is connected directly to the onboard PMIC. RTC accuracy is +/-10 seconds/day.

Table 3-2. PMIC_BBATT Pin Descriptions

Pin Name	Description	Direction	Pin Type
PMIC_BBATT	**Real Time Clock:** Optionally used to provide back-up power for the RTC in the Power Sequencer. Connects to a Lithium Cell or similar power source. The cell sources power for the RTC when system is disconnected from power. Power consumption on battery backup operation is 12 to 50µA.	Input	1.85V to 5.5V

3.7 Power Sequencing

The Jetson T5000 module and the product carrier board must be power sequenced properly to avoid potential damage to components on either the module or the carrier board system. The module is powered before the main carrier board circuits. Refer to the *Jetson Thor Series Modules Design Guide* for system level details on the application of power, power-up sequencing, power-down sequencing, and power monitoring.

3.8 Power States

The Jetson T5000 operates in three main power modes: OFF, ON, and Deep Sleep (SC7). The module transitions between these states are based on various events from hardware or software. The figure below shows the transitions between these three states.

Figure 3-1. Power State Transition Diagram

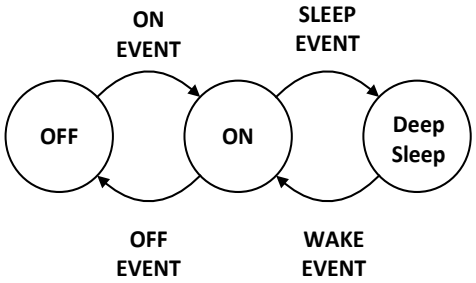

3.8.1 ON State

The ON power state is entered from either OFF or Deep Sleep states. In this state, the Jetson T5000 module is fully functional and operates normally. An ON event must occur for a transition between OFF and ON states. The VDDIN_PWR_BAD_N control is the carrier board indication to the Jetson T5000 that the VIN power is good. The carrier board should assert this high only when VIN has reached its required voltage level and is stable. This prevents Jetson T5000 from powering up until the VIN power is stable.

3.8.2 OFF State

The OFF state is the default state when the system is not powered. It can only be entered from the ON state, through an OFF event. OFF events are listed in the table below.

Table 3-3. OFF State Events

Event	Details	
Hardware Shutdown	VIN power is disconnected or MODULE_POWER_ON signal going low. The internal PMIC starts the shutdown sequence.	In ON State
Software Shutdown	Software initiated shutdown	ON state, software operational
Thermal Shutdown	If the internal temperature of the Jetson T5000 module reaches an unsafe temperature, the hardware is designed to initiate a shutdown.	Any power state

Note: Hardware shutdown, Software shutdown, and Thermal shutdown all assert MODULE_SHDN_N low. System on Module does not initiate power supply shutdown sequence until MODULE_POWER_ON is deasserted.

3.8.3 Deep Sleep (SC7) State

The SC7 state can only be entered from the ON state. This state allows the module to quickly resume to an operational state without performing a full boot sequence. In SC7 the module operates only with enough circuitry powered to allow the device to resume

and re-enter the ON state. During this state the output signals from the module are maintained at their logic level prior to entering the state (i.e., they do not change to a 0V level). To exit the SC7 state a WAKE event must occur; WAKE events can occur from within the module or from external devices through various pins on the module connector.

Table 3-4. Deep Sleep and WAKE Events

Event	Details
Thermal Condition	If the module internal temperature exceeds programmed hot and cold limits the system is forced to wake up, so it can report and take appropriate action (shut down for example).
USB VBUS detection	If VBUS is applied to the system (USB cable attached) then the device can be configured to Wake and enumerate.
Module connector Interface WAKE signals	Programmable signals on the module connector.

Note: Refer to *Jetson Thor Series Modules Pinmux Table* for WAKE capable pins.

3.9 Thermal and Power Monitoring

The Jetson T5000 is designed to operate under various workloads and environmental conditions. It has been designed so that an active or passive heat sink solution can be attached. The module contains various methods through hardware and software to limit the internal temperature to within operating limits. See the *Jetson Thor Series Modules Thermal Design Guide* for more details.

Chapter 4. Interface Descriptions

The following sections outline the interfaces available on the Jetson T5000 module. Refer to the *Jetson Thor Series Modules Design Guide* (DG-12084-001) for details of the module pins used to interact with and control each interface.

4.1 MIPI Camera Serial Interface (CSI)

Standard
MIPI® CSI 3.0 Receiver specification
MIPI C-PHY℠ 2.1
MIPI D-PHY℠ 2.1 Physical Layer specification

The NVIDIA Camera Serial Interface (NVCSI) works with the Video Input (VI) unit to capture an image from a sensor, where NVCSI is a source of pixel data to VI. NVCSI works in streaming mode while VI captures the required frames using a single-shot mode of operation. All sync point generation for software is handled at VI; the delay between NVCSI and VI is negligible in software terms. NVCSI does not have a direct memory port, instead it sends the pixel data to memory through the VI.

Fifth-generation NVIDIA camera solution (NVCSI 2.0, VI 6.0, and ISP 2.x) provides a combination host that supports enhanced MIPI D-PHY (with lane deskew support) physical layer options in two 4-lane or four 2-lane configurations; or combinations of these. Thor can support up to 16 virtual channels (VC) per ISP and supports data type interleaving.

- Virtual Channel Interleaving: VCs are defined in the CSI-2 specification and are useful when supporting multiple camera sensors. With the VC capability, a one-pixel parser (PP) can de-interleave up to 16 image streams.
- Data Type Interleaving: In HDR line-by-line mode, the sensor can output long/short exposure lines using the same VC and a different programmable data type (DT).
- Frequency Target: The parallel pixel processing rate, measured in pixels-per-clock (PPC), is increased to allow higher throughput and lower clock speeds. To support higher bandwidth without increasing the operating frequency, the host processes multiple pixels in one clock. NVCSI is capable of processing four PPCs when bits-per-pixel (BPP) is greater than 16, and eight PPC when BPP is less than or equal to 16.

- With the new streaming mode in NVCSI, one PP can handle all traffic (embedded data and image data) from one camera device, including 16 VCs.

Features:

- Supports the MIPI D-PHY v2.1 physical layer option
- MIPI D-PHY supports up to 2.5 Gbits/sec per pair, for an aggregate bandwidth of 40 Gbps from eight pairs
- Supports MIPI C-PHY v2.1
- MIPI C-PHY supports up to 4.5 Gsps (10.25 Gbps) per trio link, for a total of 164 Gbps
- Based on MIPI CSI-2 v3.0 protocol stack
- Includes six-pixel parsers (PP)
- Supports up to 16 virtual channels per active PP
- Supported input data formats:
  - RGB: RGB888, RGB666, RGB565, RGB555, RGB444
  - YUV: YUV422-8b, YUV420-8b (legacy), YUV420-8b
  - RAW: RAW6, RAW7, RAW8, RAW10, RAW12, RAW14
  - DPCM (predictor 1): 14-10-14, 14-8-14, 12-8-12, 12-7-12, 12-6-12, 12-10-12, 10-8-10, 10-7-10, 10-6-10 (Predictor 2 not supported)
- Data Type Interleave support

Note: Qualified maximum speed for C-PHY is 2.5 Gsps. Contact your NVIDIA support team if your application needs higher speed.

4.2 Camera Over Ethernet (CoE)

Transferring image data over Ethernet connection allows the reuse of existing wiring and bridge infrastructure available in a system. Instead of point-to-point communication for each camera module and SoC, a shared Ethernet network can be used, reducing the amount of cables.

4.3 Universal Serial Bus (USB)

Standard	Notes
Universal Serial Bus Specification Revision 3.2 Gen1 and Gen2	
Universal Serial Bus Specification Revision 2.0	- Modes: Host and Device (only USB 2.0 port USB0 supports RCM, Host, Device Mode. All other ports are Host only) - Speeds: Low, Full, and High - USB Battery Charging 1.2 Specifications
Enhanced Host Controller Interface Specification for Universal Serial Bus revision 1.0	

An xHCI/Device controller (named XUSB) supports the xHCI programming model for scheduling transactions and interface managements as a host that natively supports USB 3.2, USB 2.0, and USB 1.1 transactions with its USB 3.2 and USB 2.0 interfaces. The XUSB controller supports USB 2.0 L1 and L2 (suspend) link power management and USB 3.2 U1, U2, and U3 (suspend) link power managements. The XUSB controller supports remote wakeup, wake on connect, wake on disconnect, and wake on overcurrent in all power states, including sleep mode.

4.3.1 USB 2.0 Operation

Each USB 2.0 port (4x) operates in USB 2.0 high-speed mode when connecting directly to a USB 2.0 peripheral and operates in USB 1.1 full- and low-speed modes when connecting directly to a USB 1.1 peripheral. When operating in High-Speed mode, each USB 2.0 port is allocated with one High-Speed unit bandwidth. Approximately a 480 Mb/s bandwidth is allocated to each USB 2.0 port. All USB 2.0 ports operating in full- or low-speed modes share one full- and low-speed bus instance, which means 12 Mb/s theoretical bandwidth is distributed across these ports.

4.3.2 USB 3.2 Operation

In host mode, the USB3.2 host controller supports Gen2 Super Speed+, 10 Gbps transfer rates. In device mode, the USB3.2 controller supports Gen1 Super Speed.

Note: There is an internal USB 3.2 hub for ports 0 and 1. The hub supports 10Gbps bandwidth which would be shared between the two ports.

4.4 PCI Express (PCIe)

Standard	Notes
PCI Express Base Specification Revision 5.0	Thor meets the timing requirements for the Gen5 (32 Gbps per lane) data rates. Refer to the PCIe Base Specification for complete interface timing details. The PCIe controllers in Thor support Common Clock as well as Separate Reference No Spread (SRNS) and Separate Reference Independent Spread (SRIS) clocking architectures. Refer to the *Jetson Thor Series Modules Design Guide* for details. Although NVIDIA validates Thor design complies with the PCIe specification, PCIe software support may be limited. Specific PCIe use cases should be discussed with your NVIDIA representative.

The Jetson T5000 module integrates five PCIe controllers supporting:

- Connections to four interfaces, 1x1 + 1x1 + 1x2 + 1x8 (or 1x4).
- x1, x2, x8 (or x4) upstream and downstream AXI interfaces that serve as the control path from the Jetson T5000 to the external PCIe device.
- Gen5 (32 Gbps/s) supported on all controllers/lanes.

- Four PCIe controllers, seven lanes for a total of 144 GT/s.
- Controller #1 operates in x1 mode only (supports Root Port only)
- Controller #2 operates in x1 mode only (supports Root Port or Endpoint modes)
- Controller #3 can operate in x1, x2 mode (supports Root Port only)
- Controller #4 can operate in x1, x2, x4, x8 mode (supports Root Port or Endpoint modes)
- Controller #5 is available if Controller #4 is not used. In this case, Controller #5 can operate in x1, x2, x4 mode (supports Root Port or Endpoint modes)

4.5 Serial Peripheral Interface (SPI) Controller

The SPI controller allows a duplex, synchronous, serial communication between the controller and external peripheral devices. It consists of four signals:

- CS_N (Chip select)
- SCK (clock)
- MOSI (Initiator data out and Target data in)
- MISO (Initiator data in and Target data out)

The data is transferred on MISO or MOSI based on the data transfer direction on every SCK edge. The receiver always receives the data on the other edge of SCK.

Features of the SPI controller include:

- Initiator and target functionality.
  - Initiator: supports all modes in Table 4-1.
  - Target: Mode 1 and Mode 3 are supported for Tx, all Modes are supported for Rx
- Independent Rx FIFO and Tx FIFO.
- Software-controlled bit-length supports packet sizes of 4-bits to 32-bits.
- Packed mode support for bit-length of three (4-bit packet size), seven (8-bit packet size), 15 (16-bit packet size), and 31 (32-bit packet size).
- CSx_N can be selected to be controlled by software, or it can be generated automatically by the hardware on packet boundaries.
- Simultaneous receive and transmit supported.
- SPI1 and SPI3 support two chip-selects.
- Frequencies supported:
  - SPI Initiator: Operating frequency up to 67.5 MHz (1x or 2x load)
  - SPI Target: Operating frequency up to 45 MHz

Table 4-1. SPI Mode Descriptions

SPI Mode	Clock Polarity	Clock Phase	SCK Inactive State	Data Latch In	Data Latch Out
0	0	0	Low	Latched IN on the positive edge of clock	Latched OUT on the negative edge of clock
1	0	1	Low	Latched IN on the negative edge of clock	Latched OUT on the positive edge of clock
2	1	0	High	Latched IN on the negative edge of clock	Latched OUT on the positive edge of clock
3	1	1	High	Latched IN on the positive edge of clock	Latched OUT on the negative edge of clock

Figure 4-1. SPI Initiator Timing

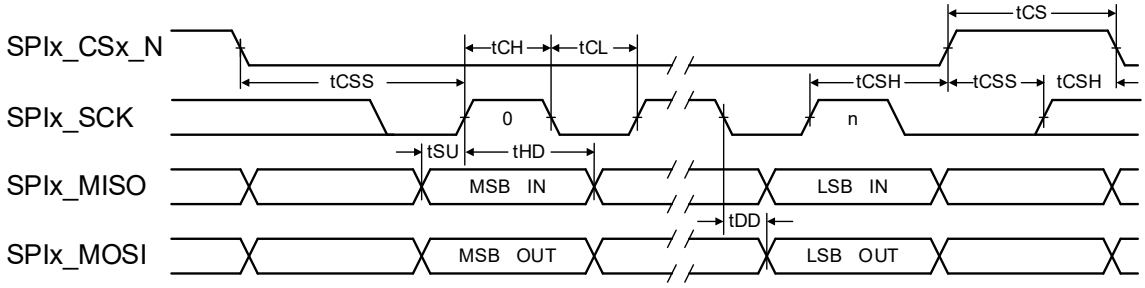

Table 4-2. SPI Initiator Timing Parameters

Symbol	Parameter	Min	Max	Unit
Fsck	SCK Clock Frequency	1	67	MHz
tSCP	SCK Period	1000 * 1/Fsck(max)	1000 * 1/Fsck(min)	ns
tCH	SCK high time	45% * tSCP	55% * tSCP	ns
tCL	SCK low time	45% * tSCP	55% * tSCP	ns
tCRT	SCK Rise time (slew rate)	0.1		V/ns
tCFT	SCK Fall time (slew rate)	0.1		V/ns
tSU	Data setup time (MISO)	2		ns
tHD	Data hold time (MISO)	3		ns
tDD	Active Clock edge to MOSI data Output Valid		6	ns
tCSS	CSx_N setup time	2		ns
tCSH	CSx_N hold time	3		ns
tCS	CSx_N high time	10		ns

Figure 4-2. SPI Target Timing

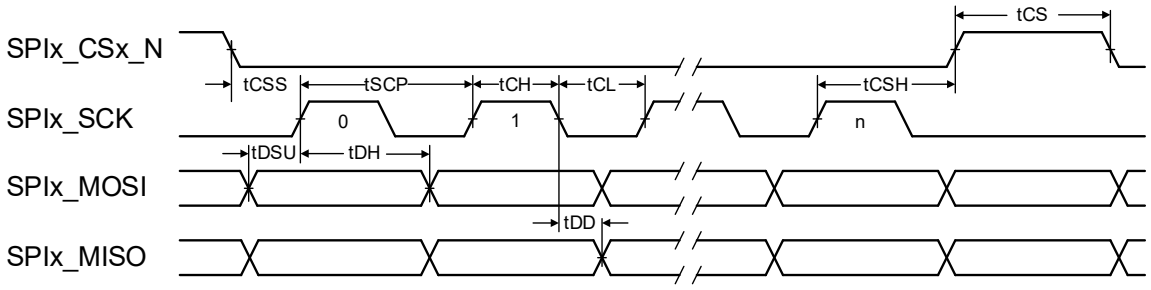

Table 4-3. SPI Target Timing Parameters

Symbol	Parameter	Min	Max	Unit
Fsck	SCK Clock Frequency	1	44.66	MHz
tSCP	SCK Period	20	1111	ns
tCH	SCK high time	45% * tSCP	55% * tSCP	ns
tCL	SCK low time	45% * tSCP	55% * tSCP	ns
tCSS	CSx_N setup time	1 * tSCP		ns
tCSH	CSx_N hold time	1 * tSCP		ns
tCS	CSx_N high time	1 * tSCP		ns
tSU	Data setup time (MOSI)	4		ns
tHD	Data hold time (MOSI)	2		ns
tDD	Active Clock edge to MISO data Output Valid	2.5	17	ns
tHO	MISO Output Hold Time	2		ns

4.6 Quad Serial Peripheral Interface (QSPI) Controllers

Jetson T5000 has two Quad Serial Peripheral Interface (QSPI) controllers: one for boot storage and one for secure key storage. Both QSPI controllers use the Serial Peripheral Interface (SPI) protocol over one or two data lanes, and the boot port additionally can operate over four data lanes. Both QSPI controllers work only as an Initiator on the bus and have independent transmit and receive FIFOs of 64 × 32 bits each.

Features:

- Initiator functionality only
- Boot port supports both SDR and DDR; Secure Key port supports only SDR
- Only Mode 0 (for both SDR and DDR) is supported for both transmit and receive operations
- Both ports support SPI mode (×1 Mode) and Dual mode (x2 mode); Boot port additionally supports Quad mode (×4 mode)
- Independent Rx and Tx FIFOs
- FIFO depth of 64 × 32 bits
- Byte-aligned bit-length in packed and unpacked mode. (BIT_LEN =7/15/31)
- PIO or DMA Mode depending on packet size
- Packed/Unpacked mode support for bit lengths of 7 (8-bit packet size), 15 (16-bit packet size), and 31 (32-bit packet size)
- In packed mode, it is expected that a minimum of one word will be transferred. Software can achieve less than one word transfer using unpacked mode.
- Least Significant Bit, Least Significant Byte First (Endian-ness),
- Only one CS support with programmable CS polarity for chip select, hardware/software CS control
- DMA support (AXI based)
- Support for combined sequences required for WRITE or READ transactions
- Hardware calibration support in QSPI controller
- Supports 1.8V QSPI flash devices only

4.7 Universal Asynchronous Receiver/Transmitter (UART)

The UART controller provides serial data synchronization and data conversion (parallel-to-serial and serial-to-parallel) for both receiver and transmitter sections. Synchronization for serial data stream is accomplished by adding start and stop bits to the transmit data to form a data character. Data integrity is accomplished by attaching a parity bit to the data character. The parity bit can be checked by the receiver for any transmission bit errors.

Note: The UART receiver input has low baud rate tolerance in 1-stop bit mode. External devices must use two stop bits. In 1-stop bit mode, the UART receiver can lose sync between the receiver and the external transmitter resulting in data errors/corruption. In 2-stop bit mode, the extra stop bit allows the UART receiver logic to align properly with the UART transmitter.

Features:

- 4x UART Interface
- Synchronization for the serial data stream with start and stop bits to transmit data and form a data character.
- Supports both 16450- and 16550-compatible modes. Default mode is 16450.
- Device clock up to 68 MHz, baud rate of 4.25 Mbits/second.
- Data integrity by attaching parity bit to the data character.
- Support for word lengths from five to eight bits, an optional parity bit and one or two stop bits.
- Support for modem control inputs.
- DMA capability for both TX and RX.
- 8-bit x 36 deep TX FIFO
- 11-bit x 36 deep RX FIFO. Three bits of 11 bits per entry will log the RX errors in FIFO mode (break, framing, and parity errors as bits 10, 9, 8 of FIFO entry)
- Auto sense baud detection.
- Timeout interrupts to indicate if the incoming stream stopped.
- Priority interrupts mechanism.
- Flow control support on RTS and CTS.
- Internal loopback
- SIR encoding/decoding (3/16 or 4/16 baud pulse widths to transmit bit zero).

4.8 Controller Area Network (CAN)

Standard	Notes
ISO/DIS 16845-2	CAN conformance test
ISO 11898-1:2015	Data link layer and physical signaling; CAN FD Frame formats
ISO 11898-4:2004	Time-triggered communication

The Controller Area Network (CAN) is a vehicular bus standard for communication between microcontrollers and devices within the vehicle. The CAN bus is a multi-Initiator serial bus for connecting multiple nodes within a vehicle using a message-based protocol. Thor FSI supports connectivity to four CAN networks.

Features:

- CAN protocol version 2.0A, version 2.0B and ISO 11898-1:2006/11898-1:2015
  - Support ISO 11898-1:2006 FD format and BOSCH FD format
  - Dual clock source, enabling FM-PLL designs
  - 16, 32, 64 or 128 Message Objects (configurable)
  - Each Message Object has its own Identifier Mask
  - Programmable FIFO mode
  - Programmable loop-back mode for self-test

- Parity check for Message RAM (optional)
  - Maskable interrupt, two interrupt lines
  - Power-down support
- Supports TTCAN
  - TTCAN Level 0, 1, and 2
  - Time Mark Interrupts
  - Stop Watch
  - Watchdog Timer
  - Synchronization to external events
- Supported bit rates up to 8 Mbps

4.9 Inter-Chip Communication (I2C)

Standard	Notes
NXP inter-IC-bus (I2C) specification	https://i2c.info/i2c-bus-specification

This general purpose I2C controller allows system expansion for I2C-based devices as defined in the NXP inter-IC-bus (I2C) specification. The I2C bus supports serial device communications to multiple devices. (13x I2C) The I2C controller handles clock source negotiation, speed negotiation for standard and fast devices, 7-bit slave address support according to the I2C protocol and supports master and slave modes of operation.

The I2C controller supports the following operating modes:

- Initiator – Standard-mode (up to 100 Kbit/s), Fast-mode (up to 400 Kbit/s), Fast-mode plus (Fm+, up to 1 Mbit/s).
- Target – Standard-mode (up to 100 Kbit/s), Fast-mode (up to 400 Kbit/s), Fast-mode plus (Fm+, up to 1 Mbit/s).

4.10 Inter-IC Sound (I2S)

Standard
Inter-IC Sound (I2S) specification

The I2S controller transports streaming audio data between system memory and an audio codec. The I2S controller supports I2S format, left-justified mode format, right-justified mode format, and DSP mode format, as defined in the Philips inter-IC-sound (I2S) bus specification.

The I2S and PCM (master and slave modes) interfaces support clock rates up to 24.5760 MHz.

The I2S controller supports point-to-point serial interfaces for the I2S digital audio streams. I2S-compatible products, such as compact disc players, digital audio tape

devices, digital sound processors, and those with digital TV sound may be directly connected to the I2S controller.

The controller also supports the PCM and telephony mode of data-transfer. Pulse-Code-Modulation (PCM) is a standard method used to digitize audio (particularly voice) patterns for transmission over digital communication channels. The Telephony mode is used to transmit and receive data to and from an external mono CODEC in a slot-based scheme of time-division multiplexing (TDM). The I2S controller supports Bidirectional audio streams and can operate in half-duplex or full-duplex mode.

Features:

- Basic I2S modes to be supported (I2S, RJM, LJM, and DSP) in both master and slave modes.
- PCM mode with short (one bit-clock wide) and long-fsync (two bit-clock wide) in both master and slave modes.
- NW-mode with independent slot-selection for both transmit and receive.
- TDM mode with flexibility in number of slots and slot(s) selection.
- Capability to drive-out a high-z outside the prescribed slot for transmission.
- Flow control for the external input/output stream.

4.11 Ethernet

Standard	Notes
Gigabit Ethernet (GbE)	IEEE 802.3ab

The Jetson T5000 module provides an Ethernet interface to support Ethernet functionality. The Ethernet PHY, magnetics and RJ45 connector are not included on the module and must be implemented externally to the module.

4.11.1 Multi-Gigabit Ethernet (MGBE)

Jetson T5000 has four integrated Multi-Gigabit Ethernet (MGBE) controllers which can support up to 100 Gbps of total bandwidth.

Each MGBE controller can independently operate up to 25 Gbps throughput mode, enabling Thor to transmit and receive data over Ethernet in compliance with IEEE 802.3-2015 standard.

The Thor MGBE controllers can be connected to external devices (e.g., Ethernet PHYs and Switches) via XFI differential lanes, enabling line rates of 2.5 Gbps, 5 Gbps, and 10 Gbps, or via 25GAUI differential lanes supporting 25 Gbps line rate.

Thor MGBE controllers support the following:

- IEEE 802.3-2015 for Ethernet MAC
- Ethernet AVB which comprises of:
  - IEEE 1588-2008 Precision Timing Protocol

  - IEEE 802.1AS-2011 based Time Synchronization
  - IEEE 802.1AS-Rev/D5.0, Timing and Synchronization for Time-Sensitive Applications
  - IEEE 802.1Qav Queue for Time Sensitive Stream
- IEEE 802.1Qbv-2015, 802.1Qbu-2016, IEEE 802.3br for Time-Sensitive Networking (TSN) traffic
- IEEE 802.1Qaz-2011 (Energy Efficient Ethernet) and 802.1Qbb-2011 for Data Center Bridging (DCB) traffic
- IEEE 802.1Q VLAN Tag. (IEEE 802.1Q-2011 includes IEEE 802.1Qat + Qav)

Notes:
1. For 25 Gbps electrical specification, IEEE 802.3by-2016 is supported.
2. USXGMII can be used for up to 10 Gbps speeds (10 Gbps, 5 Gbps, 2.5 Gbps)

4.12 Audio Controllers and Interfaces

The Audio Controller transports streaming audio data between system memory and an audio codec. The controller supports I2S format, Left-justified Mode format, Right-justified Mode format, and DSP mode format, as defined in the Philips inter-IC-sound (I2S) bus specification. The timing in the following sections applies to any of these interfaces depending on whether they are configured for I2S or TDM mode.

The audio controller supports point-to-point serial interfaces for the I2S digital audio streams. I2S-compatible products, such as compact disc players, digital audio tape devices, digital sound processors, and those with digital TV sound may be directly connected to the I2S controller. The controller also supports the PCM and telephony mode of data-transfer. Pulse-Code-Modulation (PCM) is a standard method used to digitize audio (particularly voice) patterns for transmission over digital communication channels. The Telephony mode is used to transmit and receive data to and from an external mono CODEC in a slot-based scheme of time-division multiplexing. The I2S controller supports bidirectional audio streams and can operate in half-duplex or full-duplex mode.

When the DAP port operates as an I2S (Initiator and Target modes) interface, it supports clock rates up to 12.288 MHz and comply with the I2S specification.

When the DAP port operates as a TDM/PCM interface, it supports clock rates up to 49.152 MHz.

DAP Features:

- Basic I2S modes to be supported (I2S, RJM, LJM, and DSP) in both Initiator and Target modes.
- PCM mode with short (one-bit-clock wide) and long-fsync (two bit-clocks wide) in both Initiator and Target modes.
- NW-mode with independent slot-selection for both Tx and Rx

- TDM mode with flexibility in number of slots and slot(s) selection.
- Capability to drive-out a High-z outside the prescribed slot for transmission
- Flow control for the external input/output stream.

4.13 Pulse Width Modulator (PWM)

The Pulse Width Modulator (PWM) is a frequency divider with a varying pulse width. The PWM runs off a device clock programmed in the Clock and Reset controller. The source can either be the OSC clock (27 MHz) or PLLP_OUT0 (202.5 MHz). The source is first divided by 256, and then again by a 13-bit register value, to generate the PWM frequency. The duty cycle is a controlled by an 8-bit register value.

4.14 JTAG

Thor has a JTAG interface that can be used for hardware debugging.

JTAG clock can be driven up to 15MHz.

Chapter 5. Pin Definitions

The functions for each pin on the module are fixed to a single Special-Function I/O (SFIO) or software-controlled General Purpose I/O (GPIO). The module has multiple dedicated GPIOs. Each GPIO is individually configurable as Output/Input/Interrupt sources with level and edge controls. SFIO and GPIO functionality is configured using Multi-purpose I/O (MPIO) pads. Each MPIO pad consists of:

- An output driver with tristate capability, drive strength controls and push-pull mode, open-drain mode, or both.
- An input receiver with either Schmitt mode, CMOS mode, or both.
- A weak pull-up and a weak pull-down.

MPIO pads are partitioned into multiple “pad control groups” with controls being configured for the group. During normal operation, these per-pad controls are driven by the pinmux controller registers.

Refer to the *NVIDIA Jetson Thor Series Modules Design Guide* for more information on pad behavior associated with different interfaces.

5.1 Power-On Reset Behavior

Each MPIO pad has a deterministic power-on reset (PoR) state. The reset state for each pad is chosen to minimize the need for additional on-board components; for example, on-chip weak pull-ups are enabled during PoR for pads which are usually used to drive active-low chip selects eliminating the need for additional pull-up resistors.

The following list is a simplified description of the module boot process focusing on those aspects related to the MPIO pins:

- System-level hardware executes the power-up sequence. This sequence ends when system-level hardware releases SYS_RESET_N.
- The boot ROM begins executing and programs the on-chip I/O controllers to access the secondary boot device.
- The boot ROM fetches the Boot Configuration Table (BCT) and boot loader from the secondary boot device.
- If the BCT and boot loader are fetched successfully, the boot ROM transfers control to the boot loader.
- Otherwise, the boot ROM enters USB recovery mode.

5.2 Deep Sleep Behavior

Deep Sleep is an ultra-low-power standby state in which the module maintains much of its I/O state while most of the chip is powered off. During deep sleep most of the pads are put in a state called Deep Power Down (DPD). The sequence for entering DPD is same across pads.

MPIO pads can vary during deep sleep. They differ regarding:

- Input buffer behavior during deep sleep
  - Forcibly disabled OR
  - Enabled for use as a GPIO wake event, OR
  - Enabled for some other purpose (e.g., a clock request pin)
- Output buffer behavior during deep sleep
  - Maintain a static programmable (0, 1, or tristate) constant value OR
  - Capable of changing state (i.e., dynamic while the chip is still in deep sleep)
- Weak pull-up/pull-down behavior during deep sleep
  - Forcibly disabled OR
  - Can be configured
- Pads that do not enter deep sleep
  - Some of the pads whose outputs are dynamic during deep sleep are of special type and they do not enter deep sleep.

5.3 GPIO

The module incorporates support for multiple dedicated GPIOs. Each GPIO can be individually configurable as an Output, Input, or Interrupt source with level/edge controls. All pins that can support GPIO functionality have this exposed in the Pinmux.

5.4 B2B Connector Pinout

A simplified version of the 699-pin B2B connector pinout is attached to the design guide. For more details refer to the pin description spreadsheet attached to the *Jetson Thor Series Module Design Guide*.

Chapter 6. Electrical, Mechanical, and Thermal Characteristics

6.1 Electrical Specifications

6.1.1 Operating and Absolute Maximum Ratings

The parameters listed in the following table are specific to a temperature range and operating voltage. Operating the module beyond these parameters is not recommended.

> ⚠ **Warning:** Exceeding the listed conditions may damage and/or affect long-term reliability of the part. The module should never be subjected to conditions extending beyond the ratings listed in this section.

Table 6-1. Recommended Operating Conditions

Symbol	Parameter	Minimum	Typical	Maximum	Unit
VDDDC	SYS_VIN_SV	3.135	3.3	3.465	V
	SYS_VIN_MV	4.75	5	5.25	V
	SYS_VIN_HV	7	12	20	V
	PMIC_BBATT	1.85	3.3	5.5	V
TOP	Operating Temperature (T_j)	-25	See the *Jetson Thor Series Thermal Design Guide* for details.	115	°C
TSTG	Storage Temperature (ambient)	-25	Room Temp.	80	°C

Absolute maximum ratings describe stress conditions. These parameters do not set minimum and maximum operating conditions that will be tolerated over extended periods of time. If the device is exposed to these parameters for extended periods of time, performance is not guaranteed, and device reliability may be affected. It is not recommended to operate the Jetson T5000 module under these conditions.

Table 6-2. Absolute Maximum Ratings

Symbol	Parameter	Minimum	Maximum	Unit	Comments
VDDMAX	SYS_VIN_SV	-0.3	3.5	V	
	SYS_VIN_MV	-0.3	5.5	V	
	SYS_VIN_HV	-0.3	20	V	
	PMIC_BBATT	-0.3	6	V	
IDDMAX	VIN Imax (SYS_VIN_HV)		22	A	Defines worst case peak current at room temp for 9V input.
	VIN Imax (SYS_VIN_MV)		6	A	
	PMIC_BBATT	-0.3	50	uA	Real-Time-Clock back-up supply.
VM_PIN	Voltage applied to any powered I/O pin	-0.5	0.5	V	When SYS_RESET* is low, the maximum voltage applied to any I/O pin is 0.5V
		-0.5	VDD + 0.5	V	VDD + 0.5V when SYS_RESET* is high and associated I/O rail powered. I/O pins cannot be high (>0.5V) before SYS_RESET* goes high.
	DD pins configured as open drain	-0.5	VDD x 1.1	V	The pin's output-driver must be set to open-drain mode.

6.1.2 Digital Logic

Voltages less than the minimum stated value can be interpreted as an undefined state or logic level low which may result in unreliable operation. Voltages exceeding the maximum value can damage and/or adversely affect device reliability.

Table 6-3. 3.3V Capable CMOS Pin Type DC Characteristics

Symbol	Description	Minimum	Maximum	Units
VIL	Input Low Voltage	-0.5	0.25 x VDD	V
VIH	Input High Voltage	0.75 x VDD	0.5 + VDD	V
VOL	Output Low Voltage (IOL = 2mA)	-	0.125 x VDD33 0.25 x VDD18	V
VOH	Output High Voltage (IOH = -2mA)	0.75 x VDD33 0.78 x VDD18	-	V

Table 6-4. Open Drain Capable CMOS Pin Type DC Characteristics

Symbol	Description	Minimum	Maximum	Units
VIL	Input Low Voltage	-0.5	0.2 x VDD	V
VIH	Input High Voltage	0.8 x VDD	1.1 x VDD	V

Symbol	Description	Minimum	Maximum	Units
VOL	Output Low Voltage (IOL = 2mA)	-	0.2 x VDD	V
VOH	Output High Voltage (IOH = -2mA)	0.8 x VDD	-	V

Notes:

1. See the *Jetson Thor Series Modules Design Guide* for details on I2C pull-up resistors on module.
2. See the *Jetson Thor Series Modules Design Guide* and *Jetson Thor Series Modules Pinmux Table* for supported pin voltage.

Table 6-5. DPAUX Pin Type DC Characteristics

Symbol	Description	Minimum	Maximum	Units
VIL	Input Low Voltage (DPAUX mode)	-	-200m	V
VIH	Input High Voltage (DPAUX mode)	200m	-	V
VOL	Output Low Voltage (IOL = 0mA) (DPAUX mode)	-	-400m	V
VOH	Output High Voltage (IOH = 0mA) (DPAUX mode)	400m	-	V
VIL	Input Low Voltage (I2C mode)	-	0.2 x VDD	V
VIH	Input High Voltage (I2C mode)	0.8 x VDD	-	V
VOL	Output Low Voltage (IOL = 2mA) (I2C mode)	-	0.125 x VDD	V
VOH	Output High Voltage (IOH = -2mA) (i2C mode)	0.875 x VDD	-	V

Notes:

1. See the *Jetson Thor Series Modules Design Guide* for details on I2C pull-up resistors on module.
2. See the *Jetson Thor Series Modules Design Guide* and *Jetson Thor Series Modules Pinmux Table* for supported pin voltage.

6.2 Environmental and Mechanical Screening

Module performance was assessed against a series of industry standard tests designed to evaluate robustness and estimate the failure rate of an electronic assembly in the environment in which it will be used. Mean time between failures (MTBF) calculations are produced in the design phase to predict a product's future reliability in the field.

Table 6-6. Jetson T5000 Environmental Testing

Stress Test	Test Conditions	Reference Standard
Thermal Cycling (Non-Op)	-40°C to 105°C, 43min/cyc, 500 cyc	JESD22-A104
System-level Power Cycling (Op)	40°C to 113°C, 16min/cyc, 5000 cyc	OPSQA0111--SLPC
Connector Insertion w/ base board (Non-Op)	Non-Op, Mount/unmount from baseboard, 30 times, FCT/DPA	NV Standard
Temp/Humidity Biased (Biased)	Biased (Power ON mode), 85°C, 85%RH, 168hrs	ISO 16750-4 para 5.7 IEC 60068-2-78
Low/High Temp Storage (Non-Operational)	OFF, -40°C 24hrs then 85°C 48hrs	ISO 16750-4 para 5.1.1 para 5.1.2
Power Temp Cycling (Operational/Non-Operational)	-25°C to 65°C, dwell 90min at -25°C and 110min at 65°C, 8hrs/cyc, 30cyc Op last two min at -25°C, and from 20°C to the end of 65°C	ISO 16750-4 para 5.3.1
Damp Heat, Cyclic (Operational/Non-Operational)	Op first two min at 55°C, 25°C to 55°C/95%RH, ramp 3hrs, dwell 9hrs, 24hrs/cyc, total 6cyc, 144hrs	ISO 16750-4 para 5.6.2.2
Condensation Test (Non-Operational)	25°C to 80°C, 98%RH to 80%, 5hrs/cyc, 5cyc	ISO 16750-4 para 5.6.2.4
Damp Heat, Constant (Operational/Biased)	Biased 20 days + 23hrs, then Op 1hr, 40°C 85%RH, 504hrs (21 days)	ISO 16750-4 para 5.7.2
Low Temp Operation (Operational)	Op constantly, -25°C 96hrs	ISO 16750-4 para 5.1.1.2
High Temp Operation (Operational)	Op constantly, 45°C, 1500hrs	ISO 16750-4 para 5.1.2.2
Altitude Test (Operational/Non-Operational)	Op constantly: 16,000 ft, 2,000 ft/min, 16hrs dwell, OFF: 40,000 ft, 2,000 ft/min, 16hrs dwell, room temp	ASTM D6653
Mechanical Shock (Operational)	Op constantly, 50G, 6msec, half sine, 10 shocks/orientation, total 60 shocks, room temp	ISO 16750-3 para 4.2.2.2
Random Vibration (Operational/Non-Operational)	Op/Non-Op, 10 to 1,000 Hz, 3Grms, 1hr/axis 3 axes X Y Z, Room Temp	ISO 16750-3 para 4.1
Hard Boot Cycling (Operational)	1 cyc: boot up then shutdown immediately 5,000cyc at 25°C 2,500cyc at -25°C 2,500cyc at 45°C	NV Standard

Stress Test	Test Conditions	Reference Standard
MTBF / Failure Rate: 1,323,037 hours / 756 FIT	Mobile Platforms (GM) T = 35°C, UCL = 90%	Telcordia (TelC4) SR-332, ISSUE4, Calculation Methodology: Parts Count (Method 1), UCL = 90%, Quality level: II
MTBF / Failure Rate: 1,760,838 hours / 568 FIT	Stationary Platforms (GF) T = 35°C, UCL = 90%	Telcordia (TelC4) SR-332, ISSUE4, Calculation Methodology: Parts Count (Method 1), UCL = 90%, Quality level: II

6.3 Storage and Handling

Table 6-7. Typical Handling and Storage Environment

Parameter	Description
Storage temperature (ambient)1	18°C to 30°C
Storage humidity	30% to 70% RH
Storage life2	Five years from NVIDIA shipment date to customers

Note: Transportation is a limited range of time that is covered by AEC grade 3 specs (-40°C to 85°C). Longer term storage at hubs, distribution points, and warehousing where climate controls are in place should follow conditions mentioned above.

Duration based on product being packed and stored in a controlled environment without power on.

6.4 Module Drawing and Dimensions

The following are the module dimensions, tolerances, and weight for the module.

Note: All dimensions are in millimeters unless otherwise specified

- Dimensions: 87.0 mm (width) × 100.0 mm (length) × 15.29 mm (height).
- Tolerances are: .X ± 0.25, XX ± is 0.1, Angle ± is 1.
- Weight: 0.306kg ±2%

Figure 6-1. Module Outline Drawing - 3D View

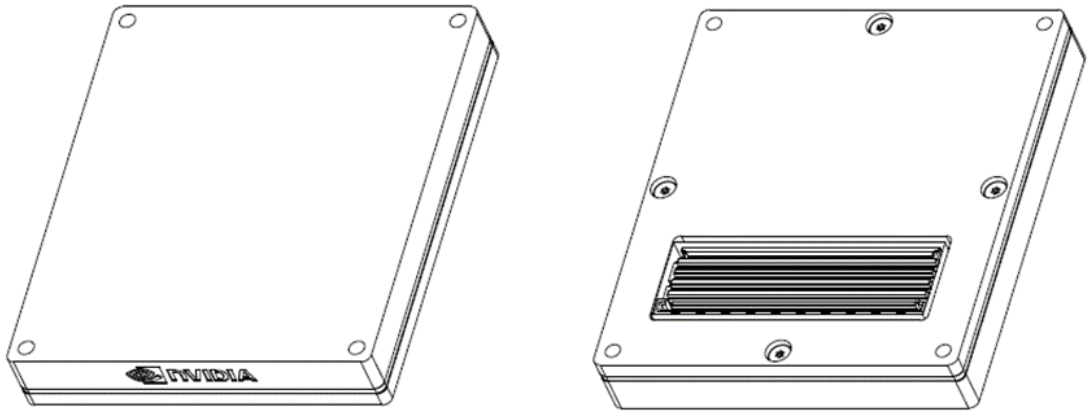

Figure 6-2. Module Mechanical Drawing - Top View

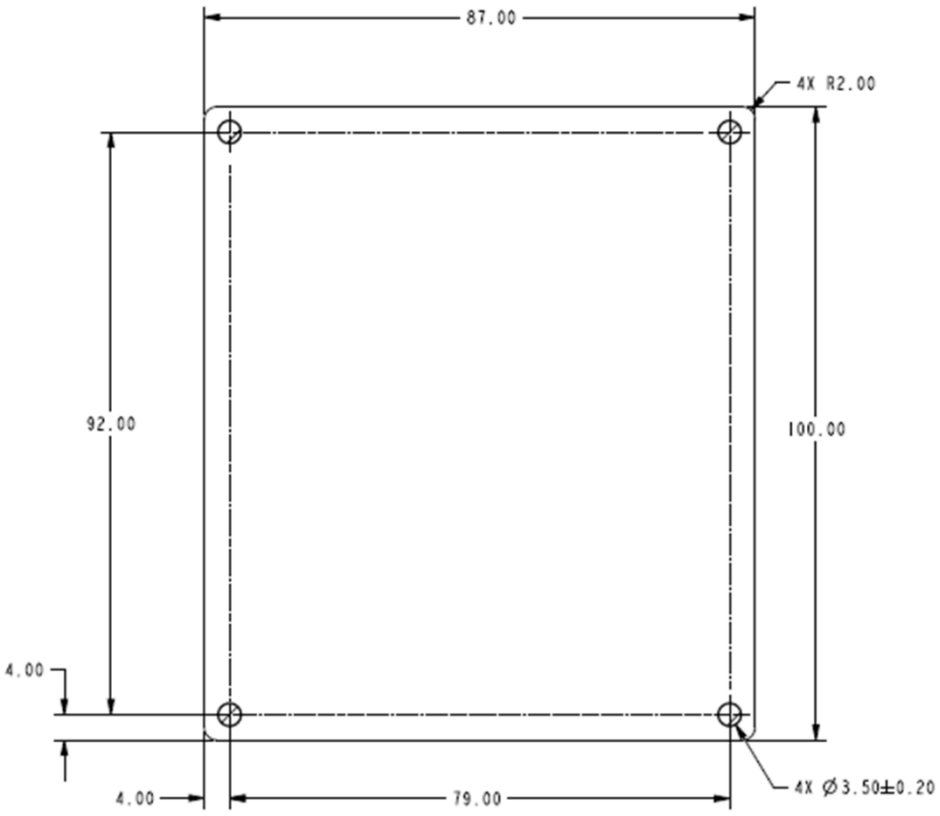

Figure 6-3. Module Mechanical Drawing - Side View

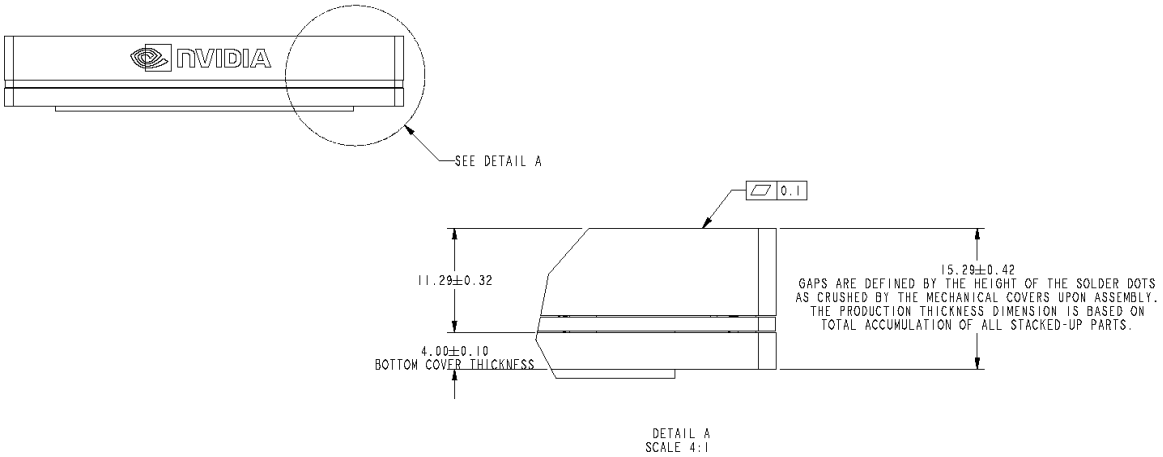

Figure 6-4. Module Mechanical Drawing - Midpoint of the Connector View

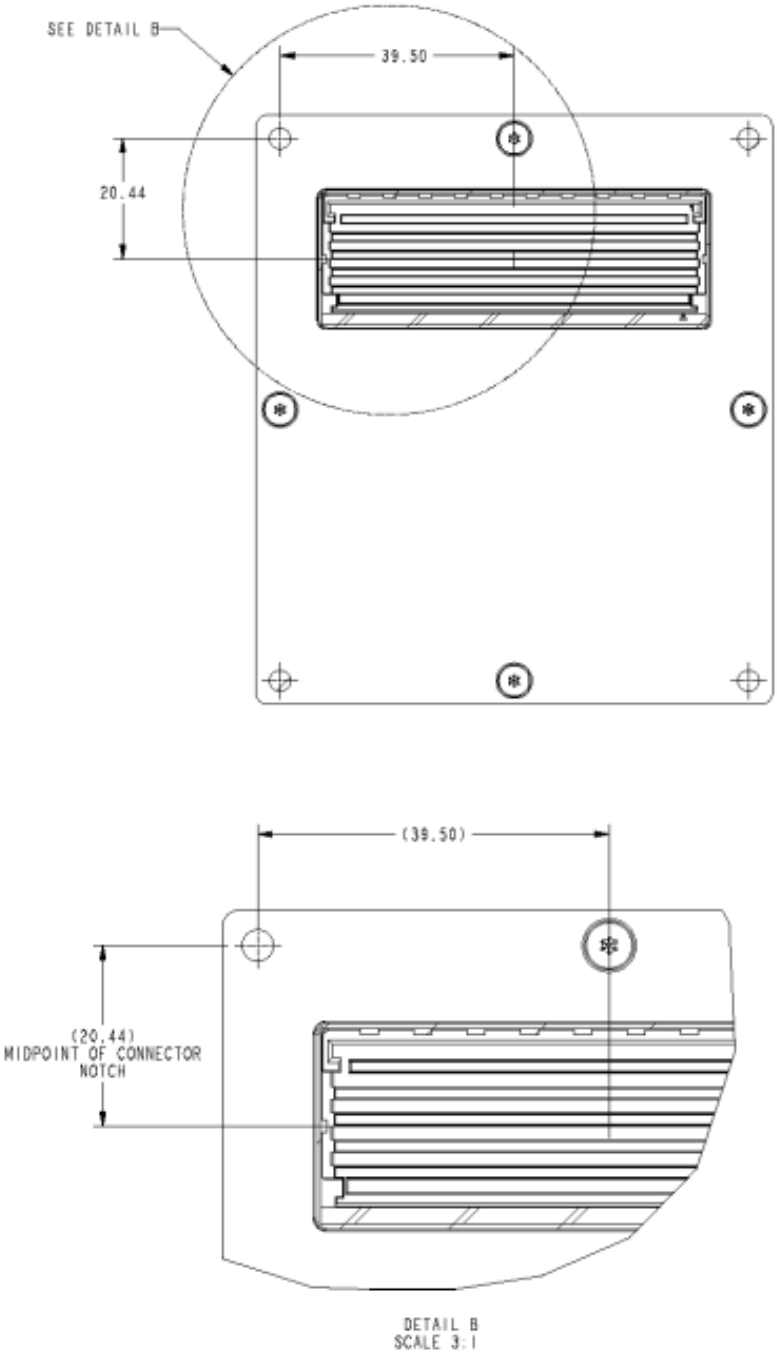

The holes labeled “A” in the following figure are used for mounting purpose to mate the NVIDIA Jetson T5000, the system motherboard, and thermal solution.

Figure 6-5. Module Mounting Holes

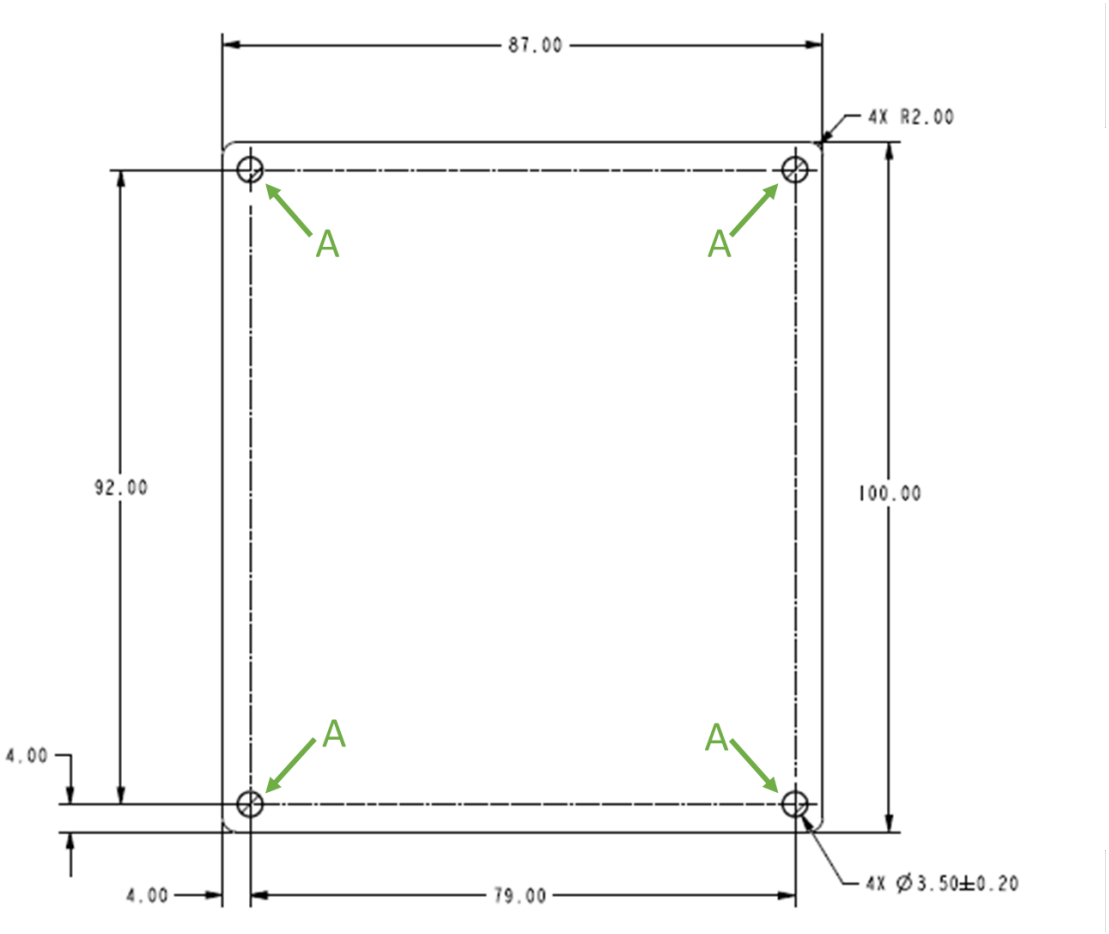

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA’s aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo, CUDA, Thor, T5000, and Jetson are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

VESA DisplayPort

DisplayPort and DisplayPort Compliance Logo, DisplayPort Compliance Logo for Dual-mode Sources, and DisplayPort Compliance Logo for Active Cables are trademarks owned by the Video Electronics Standards Association in the United States and other countries.

HDMI

HDMI, the HDMI logo, and High-Definition Multimedia Interface are trademarks or registered trademarks of HDMI Licensing LLC.

Arm

Arm, AMBA, and ARM Powered are registered trademarks of Arm Limited. Cortex, MPCore, and Mali are trademarks of Arm Limited. All other brands or product names are the property of their respective holders. ”Arm” is used to represent ARM Holdings plc; its operating company Arm Limited; and the regional subsidiaries Arm Inc.; Arm KK; Arm Korea Limited.; Arm Taiwan Limited; Arm France SAS; Arm Consulting (Shanghai) Co. Ltd.; Arm Germany GmbH; Arm Embedded Technologies Pvt. Ltd.; Arm Norway, AS, and Arm Sweden AB.

OpenCL

OpenCL is a trademark of Apple Inc. used under license to the Khronos Group Inc.

Copyright

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com