

NVIDIA Orin, ConnectX-7, HPC, Payload Profile

KEY FEATURES

- Embedded Ampere GPU with 2048 CUDA cores (5.3 TFLOPS) & 64 Gen3 Tensor cores (275 TOPS Int8)
- Embedded 12-core NVIDIA Cortex ARM64 CPU, 2.2GHz
- 64 GB LPDDR5 256-bit memory with up to 205 GB/s
- ConnectX-7, provides up to 100GbE, PCIe Gen4
- Module power: configurable from 60W - 100W

ADDITIONAL AGX ORIN FEATURES

- 2x Deep Learning Accelerator (DLA) v2 engines for inference operations
- Vision Accelerator engine for 7-way VLIW Vision Processor v2
- Dedicated programmable audio processor
- 2x HEVC (H.265) and AVC (H.264) Volta NVENC and NVDEC with up to 4K-UHD encode resolution
- CUDA® 12, OpenGL® 4.6, OpenGL ES 3.2, Vulkan™ 1.0
- Flash Storage: 64 GB eMMC 5.1 with support for ECC
- USB 2.0 (up to 480 Mb/s); optional front panel port

CONNECTIVITY / SYSTEM MANAGEMENT

- Storage: NVMe 1TB
- PCI Express Switch with x16/x32
- Backplane Ethernet with 40/100 GBASE-KR4 and 10GBASE-KR data planes; supports GPUDirect RDMA
- Block-level hardware encryption and the use of dedicated encryption keys per user
- Switching is offloaded from the CPU and run on the ConnectX hardware with NVIDIA ASAP² technology
- On-board IPMI controller for system management
- Windows, Linux or ARM root complex

MECHANICAL / OPEN SYSTEMS ARCHITECTURE

- High level of ruggedization:
 - Operating temperature: -40° to +85°C
 - Vibration (sine wave): 10G peak, 5 - 2000Hz
 - Shock: 40G peak
- Dimensions: 160mm x 100mm x 25.4mm
- Weight: TBA
- ANSI/VITA 48, 65 (VPX REDI, OpenVPX)
- SOSA Aligned Payload slot profile 14.6.11-0 (with P2 depopulated) or 14.6.13 0 (with P2B depopulated)

OVERVIEW

The VPX3U-ORIN-CX7-HPC module provides the data processing capability needed for demanding C5ISR applications, providing a secure compute node which provides advanced AI and HPC processing capabilities, high data transfer rates, and the cyber security features required to ensure data is being protected. This autonomous SOSA aligned module includes an NVIDIA Jetson AGX Orin, and an NVIDIA ConnectX-7 SmartNIC.

The NVIDIA Orin's embedded Ampere GPU provides the CUDA cores and Tensor cores for data processing, deep learning inference, machine vision, audio processing and video encoding/decoding. The 2048 CUDA cores provides 5.3 TFLOPS for processing, while the 64 Gen3 Tensor cores provides the underlying architecture required for an efficient inference engine which can achieve up to 275 TOPS (INT8) of deep learning inference computing.

The NVIDIA ConnectX-7 SmartNIC provides secure, high-speed network data transfer and a configurable PCIe switch. ConnectX-7 is ideal to meet the high data transfer and security requirements for C5ISR tasks. The ConnectX-7 also provides support for RDMA over Converged Ethernet (RoCE) and support for NVIDIA GPUDirect, enabling the fastest method for transferring data across the network to the GPU.

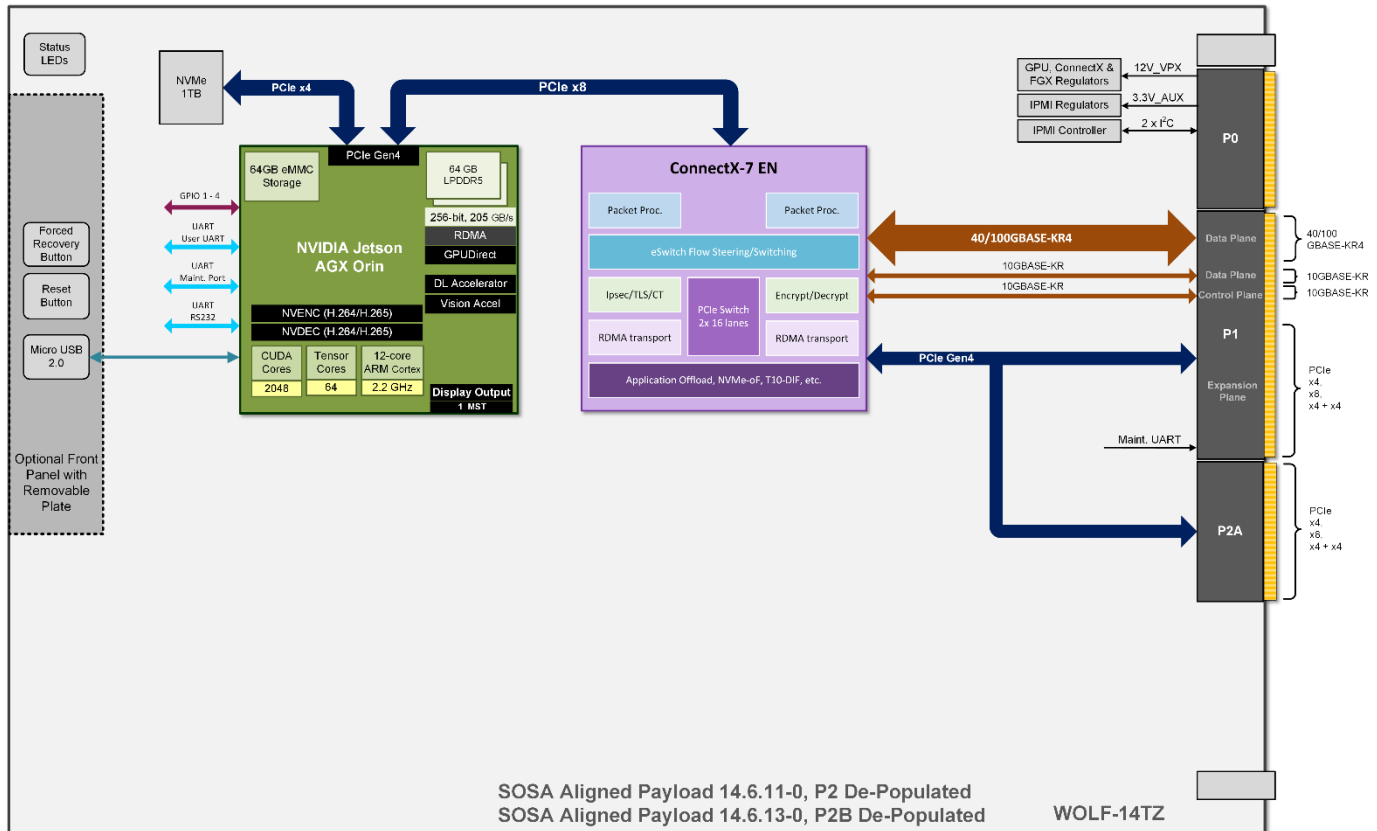
Unlocking the best performance requires the best cooling capability. WOLF's advanced cooling technology is designed to move heat using a low weight, high efficiency path from the hot chip's die to the wedgelocks.



This information is subject to change

CONNECTIVITY

The VPX3U-ORIN-CX7-HPC provides 40/100 GBASE-KR4, and PCIe Gen4 over the P1 expansion plane and optionally over the P2A expansion plane as well.



NVIDIA JETSON AGX ORIN WITH AMPERE GPU AND 12-CORE ARM CPU

Jetson AGX Orin features an embedded NVIDIA Ampere GPU with 2048 CUDA Cores and 64 Gen3 Tensor Cores, two NVIDIA deep learning accelerators, a vision accelerator, a twelve-core NVIDIA Cortex Arm CPU, and a video encoder and decoder. The NVIDIA CUDA-X accelerated computing stack and JetPack SDK support enables the Jetson AGX Orin to be a fully software-defined platform.

The Jetson AGX Orin 64GB delivers up to eight times the performance compared to the previous generation Jetson AGX Xavier. Orin can deliver up to 4 times more TFLOPS, eight times more AI TOPS, and higher memory bandwidth. Second generation NVIDIA Deep Learning Accelerators (NVDLA) can deliver up to 9 times more TOPS, while the second generation Vision Accelerator can also offer performance improvements. Connectivity has also been improved with more PCIe lanes and more available Ethernet interfaces at higher speeds. All of these performance improvements only require a modest power increase, with power settings

CONNECTX-7 ETHERNET 100GbE AND PCIe GEN4

Getting large amounts of data into and out of a module is an important system design consideration. The WOLF-12T0 module includes a ConnectX-7 SmartNIC, which provides a configurable PCIe Gen4 interface. It also provides up to 100GBASE-KR4 on the data plane, RDMA over Converged Ethernet (RoCE) with support for NVIDIA GPUDirect RDMA, and enhanced security features such as hardware-verified secure boot, hardware-accelerated cryptography, and encrypted storage.

TENSOR CORES FOR ARTIFICIAL INTELLIGENCE AND HPEC

Tensor Cores are designed to speed up the tensor / matrix computations used for deep learning neural network training and inferencing operations. NVIDIA Ampere architecture GPUs include the third-generation Tensor Core design which supports many new data types for improved performance, efficiency, and programming flexibility, including a new sparsity feature and a new Tensor Float 32 (TF32) precision mode.

NVIDIA provides CUDA-X AI and CUDA-X HPEC libraires which have been designed to work with NVIDIA Tensor Core GPUs to provide the tools needed to accelerate development of applications for AI and HPEC.

HARDWARE ACCELERATED VIDEO ENCODE / DECODE

The Ampere GPU includes the NVENC video encode (version 7.2) and NVENC decode (version 5) hardware acceleration engine. Using the Ampere GPU for video encoding provides an efficient, high quality method to achieve real time 8K and 4K encoding without burdening the system CPU. The NVIDIA Video Codec SDK provides a complete set of APIs, samples and documentation for hardware accelerated video encode and decode.

SOSA SLOT PROFILE SUPPORT

The Sensor Open Systems Architecture (SOSA) Consortium grew out of a U.S. Department of Defense (DoD) initiative to define open standard electronic architectures to ensure component interoperability, reduce costs, encourage innovation, and help to ensure a supply of needed products.

This module supports SOSA aligned Payload slot profiles. The default profile is:

- SLT3-PAY-1F1U1S1S1U1U2F1H-14.6.11-0 Payload Slot Profile, P2 depopulated
- SLT3-PAY-1F1U1S1S1U1U4F1J-14.6.13-0 Payload Slot Profile, P2B depopulated

WOLF also offers a module which supports the SOSA aligned SBC slot profile, VPX3U-XAVIER-CX7-SBC.

ORDERING CODES

The following table defines series of common order codes for the VPX3U-ORIN-CX7-HPC module. The asterisks denote characters of the part number that are defined based on common configuration options. Some common configuration options for this module are:

- Default Power Threshold
- Conformal Coatings
- PCIe Configuration Options
- Network Security Options
- Variant Locked

Ordering Number	Description
3U VPX Single Slot Configurations	
14TZ33-F***-***VPX3vAO	3U VPX, Conduction Cooled, 1", NVIDIA Orin, ConnectX-7, 40/100 GBASE-KR4, PCIe Gen4, Rear Connector P2B depopulated
14TZ33-F***-***VPX3vAO	3U VPX, Conduction Cooled, 1", NVIDIA Orin, ConnectX-7, 40/100 GBASE-KR4, PCIe Gen4, Rear Connector P2 depopulated

Contact Sales for the latest Ordering Numbers and available options

MANUFACTURING AND QUALITY ASSURANCE

WOLF designs modules to pass the following environmental standards:

- MIL-STD-810 (United States Military Standard for Environmental Engineering Considerations and Laboratory Tests)
- MIL-HDBK-217 (Reliability Prediction of Electronic Equipment)
- RTCA DO-160 (Environmental Conditions and Test Procedures for Airborne Equipment) on request

WOLF complies with the following management systems:

- AS9100D: Quality Management System - Requirements for Aviation, Space and Defense Organizations (certified)
- ISO 9001:2015: Quality management systems (certified)
- AS5553: Counterfeit Electronic Parts; Avoidance, Detection, Mitigation, and Disposition (compliant)
- NIST SP 800-171: Protecting Controlled Unclassified Information in Nonfederal Systems (compliant)

Boards are manufactured to meet the following standards:

- IPC-A-610 CLASS 3 (Acceptability of Electronic Assemblies)
- IPC 6012 CLASS 3 (Qualification and Performance Specification for Rigid Printed Boards, Class 3 for High Reliability Electronic Products)
- IPC J-STD-001 (Requirements for Soldered Electrical and Electronic Assemblies)

Caveat: integrated third party modules may not meet the same standards as WOLF manufactured modules.

