Technical White Paper

Unleashing Drone Intelligence: TDAx Processors as Next-Generation AI Companion Computers

Ruben D. Perez, Chris Tsongas

Abstract

The Texas Instruments TDAx System-on-Chip (SoC) family, developed for automotive ADAS applications, addresses the size, weight, power, and cost constraints of drone companion computers through a heterogeneous architecture in passive-cooled package options rated from -40℃ to 125℃. AEC-Q100 qualification, ISO 26262 ASIL D certification readiness, and an integrated real-time MCU safety island provide hardware reliability credentials directly applicable to drone safety certification. This paper describes the TDA4VH-Q1 architecture in the context of drone companion computer requirements and presents TI's Edge AI toolchain as a practical path to deploying production-ready AI on a supply-assured, reliability-hardened platform

Table of Contents

1 Introduction 1
2 SoC Architecture 3
2.1 Vector DSP and Matrix Math Compute Cluster 3
2.2 Vision and Motion Processing Accelerators 4
2.3 ARM® Cortex®-A72 CPU Cores 5
2.4 ARM Cortex-R5F CPU Cores 5
2.5 On-Chip Memory and High-Performance Interfaces 5
2.6 Package and Power Consumption 6
3 Machine Learning Software Development Tools 6
3.1 TI Deep Learning Tools 6
3.2 TI Edge AI Studio 6
4 TDAx Family Scalability 7
5 Conclusion 7
6 References 8

Trademarks

C7000™ and C7™ are trademarks of Texas Instruments.
ARM®, Cortex®, and TrustZone® are registered trademarks of Arm Limited.
Linux® is a registered trademark of Linus Torvalds.
All trademarks are the property of their respective owners.

1 Introduction

The Texas Instruments (TI) TDAx family of devices is well-suited for the companion computer function of drone systems to offload edge computing, artificial intelligence (AI), and computer vision tasks. This scalable family of SoC devices has been designed for automotive vision, sensor fusion, and autonomous driving in automotive driver assist systems (ADAS). While this paper focuses on airborne systems, the highlighted TDAx device capabilities can also benefit maritime applications and unmanned ground vehicles.

Modern autonomous unmanned aircraft systems (UASs) rely on a two-processor architecture that separates flight-critical functions from higher-level intelligence as shown in Figure 1-1.

Copyright © 2026 Texas Instruments Incorporated

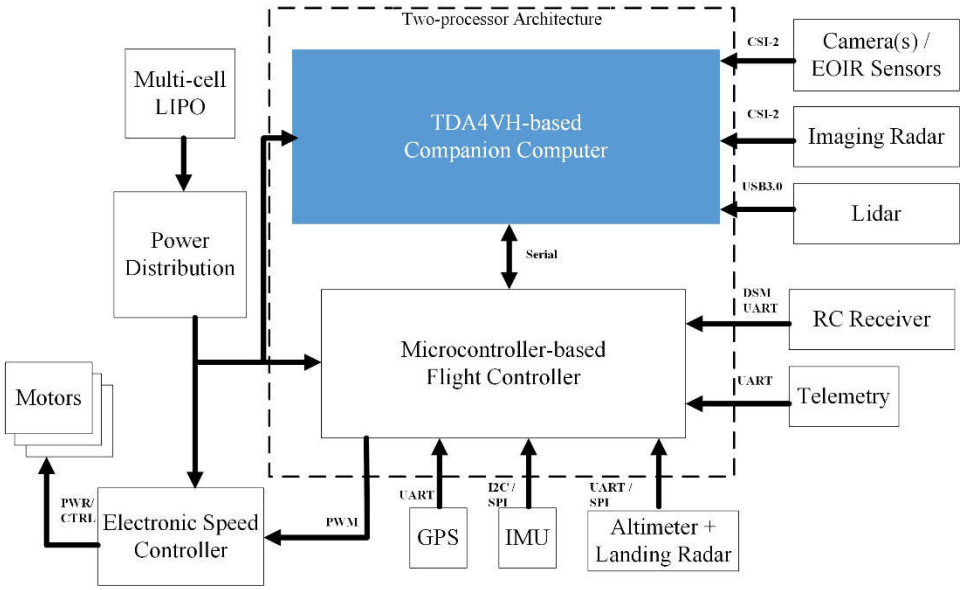


Figure 1-1. Electronics System Block Diagram of Typical Autonomous UAS

The flight controller - typically a dedicated microcontroller running real-time firmware - manages low-latency stabilization, attitude estimation, and motor control. The companion computer sits alongside it as a more powerful co-processor responsible for computationally intensive tasks that exceed the capabilities of the flight controller: sensor fusion across cameras and lidar, deep learning inference for object detection and avoidance, simultaneous localization and mapping (SLAM), mission planning, and high-bandwidth telemetry processing. The companion computer communicates with the flight controller over a serial link, translating high-level navigation decisions into velocity or position setpoint commands. As drone autonomy requirements expand from simple waypoint following to real-time scene understanding and adaptive replanning, the companion computer has become the determining factor in overall system intelligence - driving demand for processors that combine high computational throughput with the strict size, weight, power, and environmental constraints inherent to airborne platforms.

Regulatory frameworks governing commercial drone operations are simultaneously expanding the scope of permissible missions and raising the bar for system capability. FAA BVLOS waivers [7], EASA's U-Space framework [8], and Remote ID mandates [9] require drones to perform real-time detect-and-avoid, encrypted telemetry broadcast, and integration with unmanned traffic management (UTM) systems [10] - workloads that compound the AI and sensor fusion demands already placed on the companion computer. These requirements favor processors with headroom to spare: a companion computer running near its compute ceiling to process sensor data has no capacity to add compliance workloads without a platform redesign. The TDAx family's heterogeneous architecture - with dedicated acceleration for AI inference, depth processing, and general compute running concurrently - is practical for absorbing these additional regulatory workloads without displacing core autonomy functions.

Edge compute requirements similar to those of the drone companion computer have emerged in the automotive industry as car manufacturers rapidly deploy machine learning hardware in ADAS for tasks such as object detection, driver monitoring, and collision avoidance [1]. These ADAS compute platforms are highly constrained in terms of size, weight, power and cost (SWaP+C), and increasingly focused on improving energy efficiency. Ongoing investments driven by consumer demand, government safety regulations, and a competitive race to achieve Level 5 vehicle autonomy have established the market conditions in which TI TDAx processors are widely used to bring edge compute and AI acceleration to production ADAS solutions worldwide. The TDAx devices integrate general-purpose compute clusters and digital signal processing compute clusters for executing traditional signal processing and machine learning inference algorithms, in compact packages that can operate in space-constrained systems while utilizing passive cooling. These ADAS market-driven SWaP+C optimizations can now be leveraged to innovate and optimize UAS systems.

Copyright © 2026 Texas Instruments Incorporated

Several processor platforms are available to drone designers for companion computer applications, and the choice of processor has significant implications for system reliability, development cost, and long-term supply assurance. Consumer and embedded computing platforms designed primarily for cloud-edge or robotics applications typically target benign operating environments, and their industrial temperature ratings, passive thermal management capabilities, and supply continuity commitments reflect those priorities. By contrast, the TDAx family was designed from the outset to operate in the demanding thermal, vibration, and humidity environments of the automotive cabin and under-hood - conditions that closely mirror the environmental stresses encountered by airborne platforms. With a junction temperature range of -40°C to 125°C, passive cooling capability, and AEC-Q100 automotive qualification, TDAx offers a reliability advantage that consumer-grade AI accelerators lack. Unlike consumer silicon, which requires significant derating or active thermal management, TDAx avoids the associated payload and power budget penalties critical to drone performance.

2 SoC Architecture

The TDA4VH-Q1 [2] described in this paper is the most integrated member of TI's TDAx series. The TDA4VH-Q1 is built on a 16nm FinFET process and integrates a set of heterogeneous compute clusters, each purpose-built to implement specific tasks efficiently. Integrated on-chip memory and high-performance peripherals enable efficient data transfer both within the TDA4VH-Q1 SoC, as well as to and from external devices and drone subsystems.

The TDA4VH-Q1 carries the -Q1 suffix, denoting full AEC-Q100 automotive qualification [11] and support for ISO 26262 [12] functional safety development up to ASIL-D at the system level. This qualification is increasingly relevant to commercial drone designers pursuing regulatory approval for beyond visual line of sight (BVLOS) operations under FAA or EASA frameworks, where demonstrated hardware reliability and a documented safety development process are prerequisites. The TDA4VH-Q1 includes dedicated safety mechanisms such as lockstep CPU cores, error-correcting code (ECC) memory protection, and built-in self-test (BIST) capability - features absent from processors designed for consumer or general embedded markets. Rather than building from scratch, drone OEMs can accelerate their safety case development by leveraging TI's existing automotive safety documentation and functional safety manuals.

2.1 Vector DSP and Matrix Math Compute Cluster

The DSP compute cluster integrates a TI C7000 vector DSP core, coupled with a matrix multiply accelerator (MMA). For simplicity, the combination of C7000™ digital signal processor (DSP) and MMA is hereafter referred to as the C7™ neural processing unit (NPU). There are four instances of the C7 NPU compute engine available in the TDA4VH-Q1 DSP compute cluster, as shown in Figure 2-1.

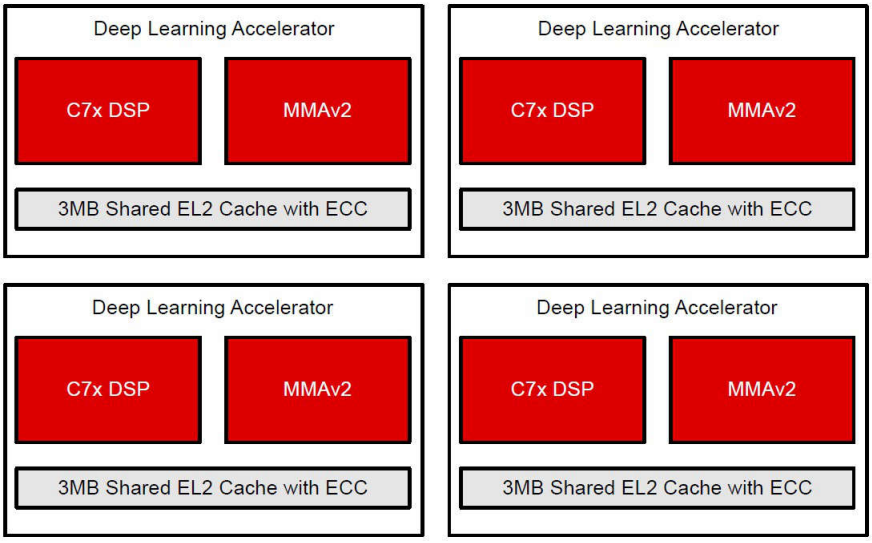


Figure 2-1. Quad C7 NPU Compute Cluster of TDA4VH-Q1

Copyright © 2026 Texas Instruments Incorporated

The 512-bit data width and efficient data-streaming architecture enable the MMA to sustain high throughput for matrix operations during machine learning inference. For instance, the MMA achieves a throughput of one 64-element inner product per clock cycle using 8-bit operands. This involves 4,096 multiply-accumulate (MAC) operations per clock cycle, totaling 8,192 operations. Because the MMA can sustain this calculation indefinitely and operates at a clock rate up to 1GHz, the MMA is capable of 8 tera operations per second (TOPS).

A separate neural network model can be computed on each of the four C7 NPU subsystems, providing a total 32 TOPS of machine learning inference performance. Alternatively, it is also possible to run a single large model spanning multiple C7 NPUs. The C7 NPU memory system is designed to provide maximal utilization of this performance; each of the four C7 NPU compute engines is tightly coupled with level 1 (L1) instruction and data caches, level 2 (L2) cache RAM, and 3Mbytes of extended L2 RAM for model storage and execution.

When not performing AI computations, the programming flexibility of the C7 DSP core enables developers to leverage the native vector DSP instructions for traditional signal processing tasks, such as Fast Fourier Transforms (FFTs), or complex matrix operations that are commonly required for SLAM computations. The C7000 Very Long Instruction Word (VLIW) architecture provides a high degree of parallelism to enable high-throughput Single Instruction Multiple Data (SIMD) operations. The Texas Instruments optimized FFT library [3] provides highly optimized kernels for computing FFTs as shown in Table 2-1.

Table 2-1. FFT Performance of a Single C7 DSP

FET Parameters		512-bit Wide C7 DSP Performance (Cycles)	
FFT Size	Number of FFTs	Int16 Data Type	32-bit Float Data Type
512	32	7,195	15,521
1024	16	9,364	17,988
2048	8	8,708	18,674
4096	4	11,281	21,945
8192	2	10,418	22,642

For example, in a range calculation for 16 data channels, a single C7 DSP can concurrently compute up to 16 1024-point floating point FFTs within 17,988 instruction cycles. When operating at the full 1GHz clock rate these sixteen FFT computations complete in under 18uS. However performing such calculations continuously incurs an additional cycle penalty of 12 to 17% for background data movements. Thus, factoring in a 16% overhead for memory transfers yields a sustained throughput is 767,000 1024-point FFTs per second.

2.2 Vision and Motion Processing Accelerators

In addition to the C7 NPU compute cluster, the TDA4VH-Q1 also includes on-chip vision processing accelerators (VPAC) and depth and motion perception accelerators (DMPAC). These accelerators handle image and motion pre-processing and are designed to work in conjunction with the deep learning algorithms for computer vision applications. The VPAC pipeline performs hardware-accelerated lens distortion correction, noise filtering, and image statistics extraction, enabling camera frames to be conditioned for AI inference without consuming C7 NPU cycles. This is particularly valuable in multi-camera drone configurations where several independent video streams must be processed simultaneously.

The DMPAC computes dense stereo depth maps and per-pixel optical flow vectors in hardware, offloading two of the most computationally demanding operations in autonomous drone navigation. Stereo depth maps provide the obstacle proximity data required for real-time collision avoidance. Dense optical flow provides a pixel-level motion estimate that can be fused with inertial measurement unit (IMU) data to generate velocity estimates for stable hover in GPS-denied environments - a capability critical for indoor inspection, tunnel survey, and contested-spectrum operations. Both outputs are produced with deterministic latency, making them practical for closed-loop feedback into the flight controller via Micro Air Vehicle Link (MAVLink) [13] or Robot Operating System 2 (ROS 2) [16].

Copyright © 2026 Texas Instruments Incorporated

2.3 ARM® Cortex®-A72 CPU Cores

There are eight ARM® Cortex®-A72 cores on TDA4VH-Q1, with clock speeds up to 2GHz. The eight cores are organized as two independent quad-core clusters, each with its own dedicated 2MB shared L2 cache, in addition to the 48KB L1 instruction cache and 32KB L1 data cache on each individual core. The two-cluster arrangement enables each cluster to be assigned to a separate operating system instance simultaneously - for example, running Linux® on one cluster and a real-time operating system such as QNX or FreeRTOS on the other - without the overhead of a full software hypervisor. Hardware virtualization support at the ARMv8-A EL2 exception level is available for more complex multi-OS configurations. ARM TrustZone® security extensions provide hardware-enforced isolation between secure and non-secure execution environments, supporting trusted execution environments and secure boot flows relevant to encrypted telemetry and remote ID compliance workloads.

This cluster is used to run the overall system application and general-purpose tasks, typically within a high-level operating system such as Linux. Machine learning acceleration can also be hosted within this cluster, providing a fallback execution path for models that have not yet been optimized for the C7 NPU, and allowing for maximum design flexibility during development. TI's Processor SDK provides pre-integrated support for Linux, QNX, and FreeRTOS on TDA4VH-Q1, with OpenVX middleware and TIDL runtime libraries enabling the A72 cluster to coordinate workloads across all compute elements on the SoC.

2.4 ARM Cortex-R5F CPU Cores

The TDA4VH-Q1 integrates eight ARM Cortex-R5F cores operating at up to 1GHz, organized across two distinct subsystems with different roles and safety characteristics. Two of these cores form a dedicated MCU safety island that operates on an isolated power domain separate from the main SoC compute enabling them to execute firmware entirely independently of the Cortex-A72 cluster, and to continue operating even if the main compute domain is powered down or experiences a fault. The remaining six Cortex-R5F cores reside in the main compute domain and are used for real-time middleware, sensor interface management, and inter-processor communication tasks that require deterministic timing but do not carry the same isolation requirements as the safety island.

In a conventional drone architecture, the flight controller function - stabilization loops, attitude estimation, and motor control - is hosted on a separate dedicated microcontroller outside the companion computer, communicating through MAVLink or DroneCAN. This separation is appropriate for most platforms and allows independent selection and qualification of the flight controller hardware. However, for designs where SWaP constraints are severe - such as fixed-wing micro-UAS, single-board nano-drone platforms, or tightly integrated payload modules - the embedded MCU capability of the TDA4VH-Q1 offers the possibility of consolidating the flight controller and companion computer functions onto a single SoC. This integration path is not available on general-purpose AI accelerators that lack dedicated real-time MCU subsystems with certified safety isolation.

2.5 On-Chip Memory and High-Performance Interfaces

The TDA4VH-Q1 SoC architecture integrates a sophisticated arrangement of tightly coupled on-chip RAM and multi-level cache systems for each compute cluster, alongside 8Mbytes of layer 3 (L3) shared RAM to maintain throughput and data sharing within the device efficiently. The SoC utilizes a shared RAM architecture across the ARM, DSP, and MMA cores; allowing data to be shared via pointer passing rather than through costly memory copies. Externally, the TDA4VH-Q1 includes four memory interface modules that are compatible with LPDDR4 memory, providing a total DDR throughput of 68 gigabytes per second (GB/s). Additionally, several high-speed interfaces, such as 10Gb/s Ethernet and multi-lane PCIe to enable flexible system dataflows.

Integration with the drone flight controller is achieved through the TDA4VH-Q1's flexible serial and networking peripherals. The device supports UART, CAN-FD, and Gigabit Ethernet interfaces, all of which are used by popular autopilot firmware. MAVLink 2 messaging over UART or UDP enables the companion computer to receive vehicle state telemetry and issue velocity, position, or attitude setpoint commands to flight controllers running PX4 [14] or ArduPilot [15]. For designs utilizing a ROS 2 architecture, the Micro XRCE-DDS [17] bridge available in PX4 v1.14 and later allows the companion computer to interact with the flight controller as a native ROS 2 node over Ethernet, providing lower latency and higher message throughput than UART-based MAVLink. Additionally, the CAN-FD interface provides a path for integration with UAVCAN/DroneCAN [18] peripheral

Copyright © 2026 Texas Instruments Incorporated

networks, enabling direct access to ESC telemetry, GPS modules, and airspeed sensors without routing through the flight controller.

2.6 Package and Power Consumption

The TI TDA4VH-Q1 is packaged in a compact 31mm × 31mm flip chip ball grid array (FCBGA) package that enables IPC class 3 PCB routing. The operating junction temperature range of –40°C to 125°C enables designers to deploy the device in designs with passive cooling designs. The total power consumption of the TDA4VH-Q1 is use-case dependent and varies with the features enabled by the application software design. However, TI provides an easy-to-use power estimation tool [4] that designers can use to quickly assess power consumption across different use cases before designing the actual system hardware. In the worst-case with all compute cores and interfaces fully utilized, the TDA4VH-Q1 consumes under 35W at 105°C junction temperature.

For the class of drone platforms most relevant to TDAx deployment - multi-rotor systems in the 3kg to 25 kg takeoff weight range used for industrial inspection, survey, logistics, and defense applications - a 10W to 20W companion computer power budget is typical. TI's power estimation tool enables designers to model realistic workloads rather than worst-case scenarios. For example, a representative configuration running two concurrent object detection models on the C7 NPU cluster, stereo depth processing on the DMPAC, and Linux application management on the Cortex-A72 cluster operates well within the 10W to 15W range, leaving thermal and power headroom for additional workloads or more aggressive duty cycling. Smaller platforms with tighter power budgets can use lower-tier members of the TDAx family to right-size compute and power consumption for the application.

3 Machine Learning Software Development Tools

The TI development tools and software development package facilitate development and the preparation of optimized firmware images for each compute cluster.

3.1 TI Deep Learning Tools

From a machine learning software development standpoint, the TI Deep Learning (TIDL) tools [5] use a bring-your-own-model (BYOM) approach, making it possible for data scientists, or embedded system designers, to reuse existing custom tool flows or open-source AI training platforms. The following industry-standard machine learning model runtime format options are available for deploying neural network models on the SoC:

- TensorFlow Lite-based inference
- ONNX Runtime-based inference
- TVM/Neo-AI-DLR based inference

The TIDL library allows each of these runtime models to be compiled for execution exclusively on C7 NPU (TIDL-RT) or with heterogeneous execution across Cortex-A72 + C7 NPU (OSRT).

3.2 TI Edge AI Studio

In addition to TIDL, TI's Edge AI Studio [6] provides a comprehensive suite of tools for the development of AI applications on TI TDAx edge processors. Edge AI Studio is available royalty-free for all TI devices. The studio includes the following tools:

- Model Composer is a fully integrated no-code solution for training and compiling models for deployment into edge AI applications. It supports BYOD, enabling the retraining of models from the TI Model Zoo with custom data. Deploy to your local development board.
- Model Analyzer enables the evaluation of accelerated deep learning inference on remotely accessed development boards. Developers can select from hundreds of optimized trained models, or use a custom model and obtain benchmarks for latency, frames-per-second processing, DDR bandwidth and accuracy on TI TDAx SoCs.
- TI Model Zoo is a TI-hosted repository with a collection of example Deep Learning Models for various Computer Vision tasks. These tasks include image classification, segmentation and object detection. The models in this repository can be executed either in host emulation mode on an X86 PC or on a TI development board.

Copyright © 2026 Texas Instruments Incorporated

The overall work flow of DNN development and deployment on TI devices is shown in Figure 3-1.

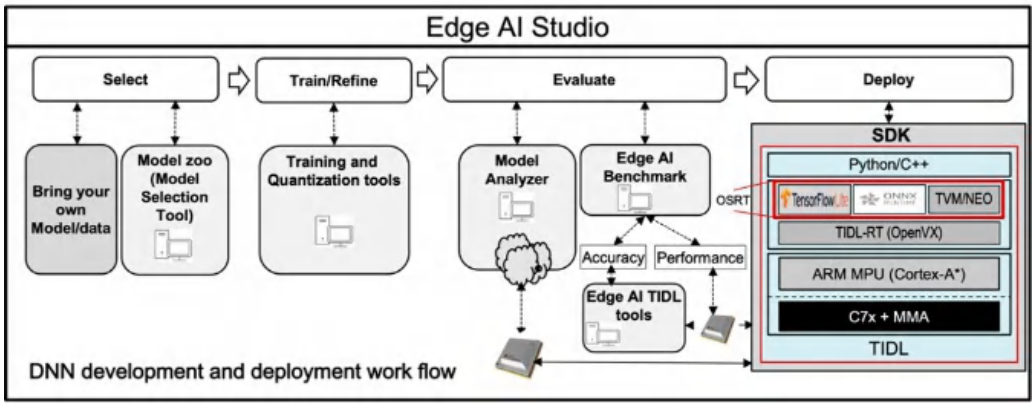


Figure 3-1. DNN Development and Deployment on TI Devices

The included training and quantization tools for popular frameworks allow data scientists to make DNNs operate efficiently on TI devices. TIDL tools also include an AI benchmark framework for developers to benchmark the accuracy and performance of edge AI model deployments on TI TDAx SoCs.

4 TDAx Family Scalability

The TDA4VH-Q1 described in this paper represents the highest-integration member of TI's TDA4x processor family. The family spans a range of production-ready options with varying compute performance and power envelopes, enabling drone designers to select the appropriate device for their platform class without rewriting application software. The Processor SDK and TIDL tool chain are common across the family, so software developed on one device can be redeployed on another with minimal modification. Furthermore, the TDA5 next-gen devices further enhance performance and IP integration. Table 4-1 summarizes the key parameters across select members of the TDAx family.

Table 4-1. TDAx Family Scalability

Processor Family Member	Application Cores (MPU)	Deep Learning (C7x DSP / MMA)	On-Chip and External Memory Interfaces
TDA4VM-Q1	2x Arm Cortex-A72	1x C7x DSP + MMA	8MBytes L3 + 1x LPDDR4
TDA4VE-Q1	2x Arm Cortex-A72	2x C7x DSP + MMA	4MBytes L3 + 2x LPDDR4
TDA4VH-Q1	8x Arm Cortex-A72	4x C7x DSP + MMA	8MBytes L3 + 4x LPDDR4
TDA54-Q1 (Next Gen)	8x Arm Cortex-A720AE	4x C7 Neural Processing Units (Up to 400TOPS)	24MBytes L3 + 4 LPDDR4/5x

This scalability provides a practical advantage for drone manufacturers developing a product family across payload classes - from a lightweight fixed-wing surveillance platform to a heavy-lift cargo UAS - where a unified AI software stack must run on processors with varying performance and power envelopes. Furthermore, the common toolchain also enables algorithm development and validation to begin on higher-performance devices before final optimization for the target platform.

5 Conclusion

The TI TDAx processor SoCs are production-proven AI compute platforms with over three years of high-volume deployment in ADAS applications, establishing a reliability and supply continuity record that purpose-built or consumer-grade alternatives cannot match. This track record goes beyond mere marketing; it demonstrates validated robustness against the thermal, vibration, and operational demands of autonomous drones, backed by a supply chain engineered to support long-term industrial and defense programs.

The architectural alignment between ADAS and drone companion computer requirements is direct: heterogeneous AI acceleration, hardware-accelerated vision and depth processing, passive thermal operation, and a wide industrial temperature range address the core SWaP+C constraints of airborne, maritime, and ground-based platforms without compromise. The TDA4x family's scalability across compute tiers, combined

Copyright © 2026 Texas Instruments Incorporated

with a common software stack built on open standards - TensorFlow Lite, ONNX Runtime, ROS 2, and MAVLink - means that drone designers inherit a mature development ecosystem rather than building from scratch.

For drone OEMs navigating an increasingly demanding regulatory environment, the TDA4VH-Q1's AEC-Q100 qualification and ISO 26262 functional safety support provide a documented foundation for BVLOS certification efforts. Combined with TI's Edge AI Studio and TIDL toolchain, the path from AI model development to deployment on a flight-ready companion computer is well-defined, well-supported, and available today.

6 References

1. P. S. Chib and P. Singh, "Recent Advancements in End-to-End Autonomous Driving Using Deep Learning: A Survey," in *IEEE Transactions on Intelligent Vehicles*, vol. 9, no. 1, pp. 103-118, Jan. 2024.
2. Texas Instruments, TDA4VH-Q1 SoC with Octal Arm® Cortex®-A72, 32 TOPS of AI, C7xDSP, and GPU for vision perception and analytics, product page
3. Texas Instruments, Processor SDK RTOS J784S4
4. Texas Instruments, J784S4, TDA4VH, TDA4AH, TDA4VP, TDA4AP, AM69 Power Estimation Tool
5. Texas Instruments, TIDL - TI Deep Learning Product, Github repository
6. Texas Instruments, CCStudio™ Edge AI Studio
7. Federal Aviation Administration, *Normalizing Unmanned Aircraft Systems Beyond Visual Line of Sight Operations*, notice of proposed rulemaking
8. European Commission, "Commission Implementing Regulation (EU) 2021/664 — Regulatory Framework for the U-Space."
9. Federal Aviation Administration, "Remote Identification of Unmanned Aircraft," 14 CFR Part 89, Final Rule, Jan. 2021
10. Federal Aviation Administration, Unmanned Aircraft System (UAS) Traffic Management (UTM) Concept of Operations, Version 2.0, concept of operations
11. Automotive Electronics Council, "AEC-Q100 Rev J: Failure Mechanism Based Stress Test Qualification for Integrated Circuits."
12. International Organization for Standardization, "ISO 26262-1:2018 — Road Vehicles: Functional Safety," 2nd ed., 2018
13. MAVLink Community / Dronecode, "MAVLink 2 Protocol Specification"
14. PX4 Development Team / Dronecode, "PX4 Autopilot User Guide." Zenodo, doi: 10.5281/zenodo.20170875
15. ArduPilot Development Team, "ArduPilot Documentation."
16. S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, "Robot Operating System 2: Design, Architecture, and Uses In The Wild," in Science Robotics, vol. 7, no. 66, 2022, doi: 10.1126/scirobotics.abm6074
17. Object Management Group, "DDS for eXtremely Resource Constrained Environments (DDS-XRCE) Specification, v1.0," 2018
18. DroneCAN Community, "DroneCAN Specification v1," Github page

Copyright © 2026 Texas Instruments Incorporated

IMPORTANT NOTICE AND DISCLAIMER

TI PROVIDES TECHNICAL AND RELIABILITY DATA (INCLUDING DATASHEETS), DESIGN RESOURCES (INCLUDING REFERENCE DESIGNS), APPLICATION OR OTHER DESIGN ADVICE, WEB TOOLS, SAFETY INFORMATION, AND OTHER RESOURCES “AS IS” AND WITH ALL FAULTS, AND DISCLAIMS ALL WARRANTIES, EXPRESS AND IMPLIED, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT OF THIRD PARTY INTELLECTUAL PROPERTY RIGHTS.

These resources are intended for skilled developers designing with TI products. You are solely responsible for (1) selecting the appropriate TI products for your application, (2) designing, validating and testing your application, and (3) ensuring your application meets applicable standards, and any other safety, security, regulatory or other requirements.

These resources are subject to change without notice. TI grants you permission to use these resources only for development of an application that uses the TI products described in the resource. Other reproduction and display of these resources is prohibited. No license is granted to any other TI intellectual property right or to any third party intellectual property right. TI disclaims responsibility for, and you fully indemnify TI and its representatives against any claims, damages, costs, losses, and liabilities arising out of your use of these resources.

TI’s products are provided subject to TI’s Terms of Sale, TI’s General Quality Guidelines, or other applicable terms available either on ti.com or provided in conjunction with such TI products. TI’s provision of these resources does not expand or otherwise alter TI’s applicable warranties or warranty disclaimers for TI products. Unless TI explicitly designates a product as custom or customer-specified, TI products are standard, catalog, general purpose devices.

TI objects to and rejects any additional or different terms you may propose.

Copyright © 2026, Texas Instruments Incorporated

Last updated 10/2025