

White Paper

The Future of Avionics: High Performance, Machine-Learned and Certified

Authors

Debra Aubrey

Technical Product Marketing Manager,
Intel Corporation

Dr. Yemaya Bordain

President of Americas,
Daedalean

Dr. Niels Haandbaek

Director of Engineering,
Daedalean

Abstract	2
Introduction	3
Embedded computing in the aerospace industry	4
Requirements at all hardware and software levels: How to prove that aviation systems work as intended and are safe	4
Certification of ML-based aviation software	5
Applying software assurance to machine learning: Additional challenges	6
• Challenge 1: Traditional software assurance methods are not applicable	6
• Challenge 2: Unprecedented computational performance required from an aerospace embedded platform.....	6
The Visual Awareness System	7
Daedalean Tensor Accelerator	9
The unique advantage of the Intel® Core™ i7 processor: Allowing certification all the way to DAL-A	9
Resource Management and AMC 20-193/CAST 32A	10
Conclusion	11

Abstract

The next revolution in aviation is on the horizon. After powered flight and the jet aircraft, today high automation, autonomy, and electric propulsion give rise to the *Jetsons*, where humans and goods are zipped across urban skies in *flying cars*. Realizing advanced air mobility requires enabling technologies that are yet to be fully demonstrated, including avionics systems that are robust enough to operate as situationally-aware pilots, both in the presence and absence of human operators. Daedalean is inventing the foundational technologies that are driving progress toward autonomy. In doing so, we are bridging the gap between the aerospace and defense (A&D) industry and the field of artificial intelligence (AI), starting with reengineering the airborne electronic systems that provide intelligence in the cockpit. We outline the path to autonomy in our [Roadmap 2028](#).

Embedded aircraft OEMs may lack knowledge and experience with machine learning (ML) and may struggle to build systems that support and incorporate ML. Safety-critical ML-based applications pose additional constraints: certifiability, low SWaP, and high computational performance. **In this whitepaper, Daedalean presents a certifiable embedded system that takes inputs from high-resolution cameras and/or other sensors (such as radar, LiDAR, etc.) and applies machine learning to, for instance, detect traffic in the airspace. Although we motivate the need for such equipment with a real-world example of a machine-learned visual awareness system, we limit the focus of this paper to computing hardware and embedded software. Therefore, details about how we assure the equipment to perform its intended function(s) and comply with applicable standards are out of scope for this paper. We foresee that availability of such equipment will enable a variety of novel use cases based on machine learning.**

Based on these use cases, Daedalean outlines the specific hardware requirements necessary to support its technology, including a minimum of one Tera Operations per Second (TOPS) and low Size, Weight and Power consumption (SWaP), and proposes a solution in the form of a reference architecture. This architecture includes the 11th Gen Intel® Core™ i7 processor (formerly known as Tiger Lake UP3), the Intel® Agilex™ F-Series FPGA, and Intel® Airworthiness Evidence Package (Intel® AEP), as well as Daedalean's Visual Awareness System, Vyper hypervisor, certifiable ML, and Situational Intelligence™ applications.

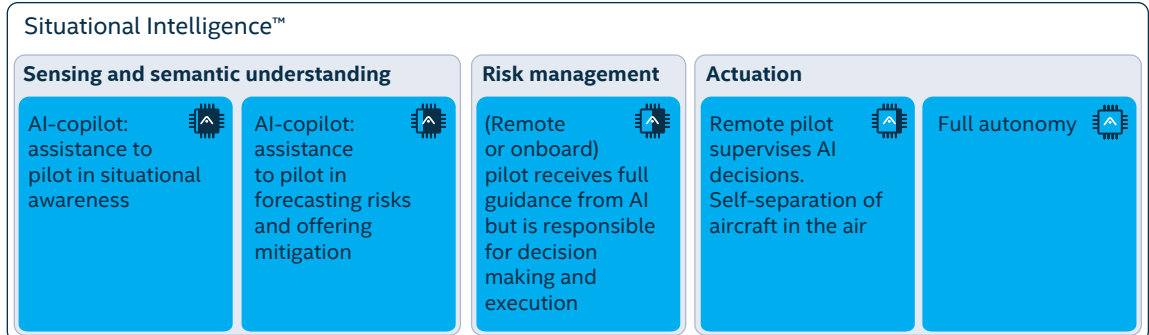
Integrating AI into A&D systems is a challenge that Daedalean has addressed. The Intel® AEP provides aircraft-embedded manufacturers with processor artifacts to support DO-254 certification up to DAL-A. By partnering with Intel, Daedalean can offer low SWaP solutions that meet both the computational and certification requirements.

This whitepaper is the first of a series aimed at embedded companies in the A&D industry looking to incorporate certifiable AI into their systems, i.e., avionics OEMs, COTS solutions providers, airframers/systems integrators. Daedalean's software seamlessly integrates with hardware architectures that meet the specifications proposed below. Likewise, Daedalean's software can integrate into existing and legacy systems. However, performance, certifiability, or both may be forfeited. Safety certification constraints and performance requirements limit the portability of Daedalean's software.

Levels of autonomy



Daedalean's roadmap for the AI and human roles



AI application integration



The stages toward autonomy and how AI/ML fit in this transition. Daedalean's Situational Intelligence™ suite provides an AI/ML co-pilot that advances its role with increased automation until it is promoted to pilot of autonomous aircraft

Introduction

Embedded computing in the aerospace industry has high demands for proof that systems work as intended without introducing unintended functions. This precondition translates to strict requirements at all levels of hardware and software. Applying machine learning in the aerospace industry presents two additional challenges: first, ensuring that machine learning algorithms work safely and efficiently, and second, that the embedded platform has enough computational power to handle the increased workloads.

Daedalean provides certifiable Situational Intelligence™ systems built on machine learning. Daedalean's co-piloting system includes vision-based traffic detection, navigation, and landing guidance functions. The first product—**PilotEye™**, a traffic avoidance system with integrated vision-based traffic detection developed in partnership with Avidyne Corporation – is poised to be the first DAL-C-certified machine-learned application. Our Situational Intelligence™ suite is available today through our **Eval Kit**, an uncertified vision-based system that can be installed on aircraft with field-approved and certified assembly for experimental purposes. In achieving the certification milestone, our AI-enhanced systems are foundational for enabling air autonomy in commercial aerospace. To that end, we are defining future avionics systems.

In order to recognize aircraft and runways in real time, Daedalean's machine learning algorithms require a minimum of one Tera (10^{12}) Operations Per Second (TOPS) per camera¹ which outputs 6 frames per second of 12 Megapixel each. Hence, to handle four camera streams, as is necessary for the full-performance visual awareness suite architecture, means processing 288 Megapixel per second and requires a minimum computational power of four TOPS.

Until recently, the venerable PowerPC architecture was the top-performing option in the aerospace industry. Initially designed in the 1990s, the last version on the market that was certifiable for use in avionics featured four cores with a maximum clock speed of 1.6 GHz. Moreover, to meet the determinism requirements of Design Assurance Level A (DAL-A), three of the four cores had to be disabled. For high-performance tensor computations, such as those used in Daedalean's applications, the only available GPUs for use in avionics were the AMD Radeon 8860 and 9170, also both End-of-Life. Together, this setup would be barely capable of supporting a single camera stream.

However, Daedalean has developed solutions that enable high-performance tensor computations and can be certified for use up to the highest level of design assurance (DAL-A). These solutions use modern hardware, such as the 11th Gen Intel® Core™ i7 processor and the Intel® Agilex™ F-series line of FPGAs. Although the first product to launch to the market was developed on FPGA alternatives, the recent availability of the Intel® Agilex™ family of FPGAs makes them a natural follow-on with significant SWaP improvements.



A conceptual view of Daedalean's Situational Intelligence™ suite, which provides sensor-fused traffic detection, navigation, and landing guidance to co-piloting systems. The display UI varies by application and OEM

The solution described in this document utilizes a vision and camera-based approach for situational awareness. The defined architecture is flexible and performant even beyond the requirements mentioned above: it is scalable to meet the computational requirements for handling a variety of other sensors, such as radar, LiDAR, and long-wave infrared cameras, for creating a more functional and versatile system for AI-enhanced situational awareness beyond visual conditions. A system that fuses these sensors is currently in development at Daedalean.

Daedalean's architecture can be integrated with non-AI-based OEM safety-critical co-applications, such as flight management systems and cockpit display systems over standard avionics-grade bus protocols..

The unique and complex requirements of certifying AI/ML systems in the A&D industry present a challenge that Daedalean has set out to overcome since its conception as a company. While commercial off-the-shelf silicon solutions may meet the computing requirements for AI and machine learning applications, they do not meet the aviation certifiability requirements for these technologies. Intel offers the Intel® Airworthiness Evidence Package for the 11th Gen Intel® Core™ i7-1186GRE processor that equips aircraft-embedded manufacturers with processor artifacts to support DO-254 certification up to DAL-A. Intel also provides essential reliability documentation supporting their Agilex™ FPGAs. By partnering with Intel and utilizing these airworthiness-enabling processing solutions, Daedalean is able to offer certifiable AI/ML technology for use in aviation applications.

¹ Determined by running Daedalean's 1-camera system on alternative HW and calculating the camera input rate and the minimum speed for processing the camera feed

Embedded computing in the aerospace industry

Normally, high performance and reliability are considered orthogonal requirements for hardware. Graphics processors used for video games and scientific computing are highly performant but not predictable (or at least it may be impossible to get the documentation to assert its predictability)—while legacy embedded hardware may be predictable but not performant. A difficult challenge for embedded developers is then to meet nonfunctional performance requirements, such as hard real-time execution bounds, memory capacity, and energy consumption. Hardware designers have gone to great lengths to hide all kinds of complexity under the hood so that medium programming effort gives the highest possible average computational throughput. In embedded solutions, extraordinary programming effort is spent on recovering determinism by undoing some of the magic performed by the hardware.

Designing embedded computing systems can be very different from designing general-purpose computing systems. It requires several extra skills: system specification, not only in the informal sense but also in creating executable models; basic architectural techniques for both hardware and software; analysis of performance and energy consumption for both hardware and software; and making trade-offs between hardware and software at all stages in the design process.

Embedded systems for avionics are even more challenging to build: they have to meet real-time and functional safety requirements for industrial use conditions. Those include the long life of parts (7-15 years) and obsolescence management; SWaP limitations; robustness against external vibrations; support for extended temperature ranges; and high performance—all in a single package.

Safety-certifiable avionics systems not only need to meet application-specific system and reliability requirements but also undergo design assurance in accordance with the established safety compliance for developing aerospace systems. In addition, the introduction of machine learning impacts the system requirements at all levels.

This white paper is the first to outline considerations for designing safety-certifiable, deterministic, high-performance machine-learned avionics systems. We discuss the aspects across hardware and software, including silicon selection, architecture, CAST 32A and AMC 20-193, neural network development, etc., and challenges when putting it together to meet regulatory requirements. We propose an architecture for a vision-based machine-learned system that meets SWaP, performance, reliability, environmental, and regulatory requirements to gain certification.

Requirements at all hardware and software levels: How to prove that aviation systems work as intended and are safe

Software assurance is a set of processes aimed at providing guarantees of the intended behavior, in particular, the absence of bugs. These processes are included in the development life cycle, from initial planning to final testing. They are not costless and not flexible. In the fields where software is critical for safety—be it aviation, health care, or industrial control—it is not acceptable to save time by skipping the “bug-free” requirement during the long design and development stage, as well as all the other procedures. Achieving software assurance may mean re-design and further development of supplier components, such as custom fast-boot BIOS, firmware development, disablement of unexposed features (such as the Intel® Converged Security Engine®, also called the “Management Engine” in other architectures), as well as

disabling thermal throttling and dynamic clock frequency scaling, i.e., Turbo Boost Technology.

Aviation systems needing certification are subject to multiple levels of requirements outlined by the following guidance:

- ARP4754A: ground-based system requirements;
- ARP 4761A: safety and derived requirements;
- DO-254: hardware requirements;
 - AC/AMC 20-152A: guidance for using COTS parts in certifiable avionics;
- DO-178C: software requirements
 - AMC 20-193: use of multicore processors;
 - CAST-32A (FAA): multicore processors.

These guidance are interrelated. ARP4754A and ARP4761 apply to the aircraft and system level whose requirements flow down to individual item levels (DO-254, DO-178C). Then implementation and verification flow back up from the item to the airframe level. With the increased use of commercial-off-the-shelf components and multi-core processors, the AC/AMC 20-152A, AMC 20-193, and CAST-32A were introduced to apply the sub-system safety requirements to modern electronics.

The “V” model is well-established, tried, and true for developing safe embedded aerospace systems. However, the introduction of machine learning has required new guidance, namely the [First Usable Guidance for Level 1 machine learning applications](#), proposed by the EASA in 2021 and heavily adapted from two publications co-authored by Daedalean and the EASA (discussed below). *The First Usable Guidance...* outlines preliminary ML requirements.

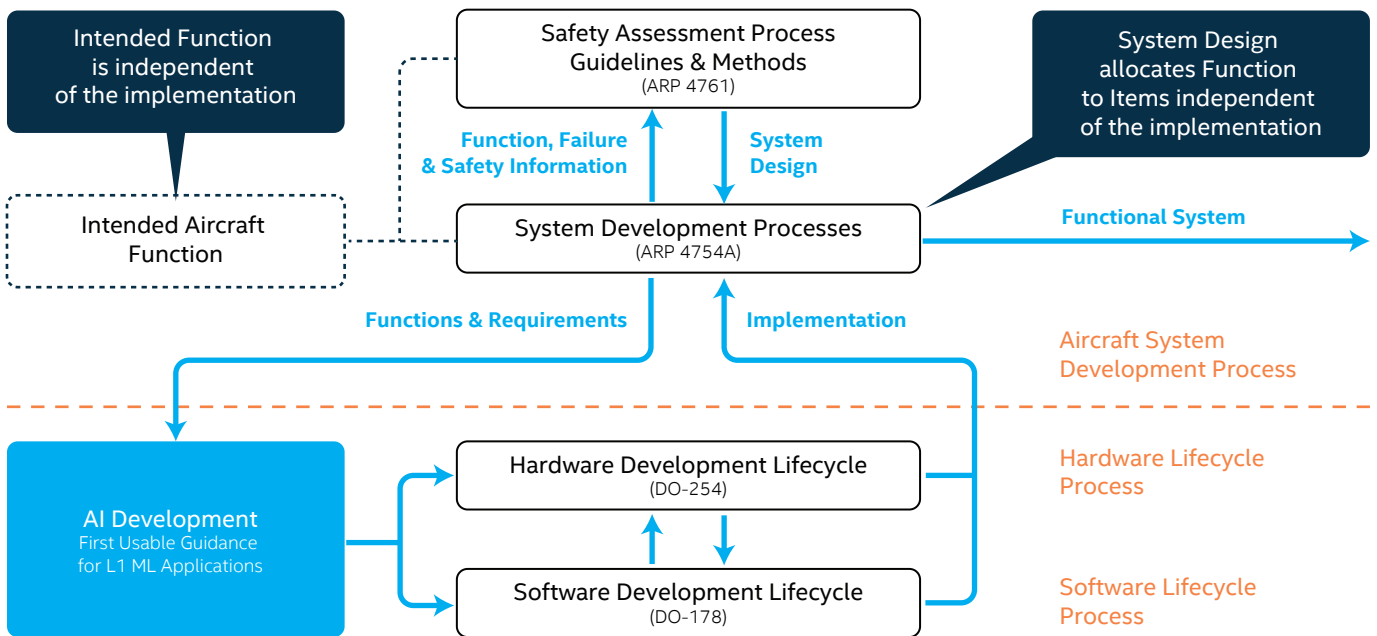
Certification of ML-based aviation software

A neural network (NN) is a piece of software designed to perform a specific function, such as analyzing data from sensors and providing conclusions (called *inferences* in ML vocabulary) based on that data. A NN does not possess the ability to make decisions or have opinions. Its behavior is determined by the parameters trained into it during the machine learning process, which can take place entirely offline (and, for safety-critical applications, *must*² be offline in the interest of being predictable).

While a neural network can generate incorrect inferences, the effect of those on the final output of the application can be mitigated through careful system design and the introduction of redundancy and, for DAL-A applications, dissimilarity. For example, by analyzing a sequence of images rather than just a single image, the system can provide more accurate answers and reduce the probability of false alerts. The level of acceptable error rate is determined by the system requirements, typically adhering to industry standards, and is used to set the bar for safety performance in aviation applications.

Machine learning-based aviation software can be certified according to aviation standards through a “learning assurance” process similar to the existing software assurance process for traditional software. This process can help to ensure that the ML-based software meets the required safety performance standards and does not produce errors more frequently in the field than is acceptable. This goes beyond merely certifying the runtime software which executes a neural network or other ML inference function to DO-178C.

Although certifying the ML inference runtime is a necessary ingredient, to further be able to assert and certify that the machine learned part actually implements the correct emerging behaviour, the neural network at the core of the system must be a frozen, deterministic program that performs a specific function, and has been trained, validated, and tested on independent annotated datasets.



Where does AI fit?

Then, by carefully decomposing the system-level requirements into requirements on the datasets and learning algorithms, Statistical Learning Theory methods can be used to obtain assurance on the behavior of the system in its operational domain. These methods lead to learning assurance. These were the subject of research conducted by EASA and the FAA in collaboration with Daedalean. Later, EASA proposed this process in their guidelines for certifica-

tion for applications based on neural networks as a possible Acceptable Means of Compliance. Daedalean’s PilotEye™ STC process is an example of this certification process in action and is the first test case for both EASA and the FAA to apply these procedures and evaluate their effectiveness.

Below we outline the challenges that inspired the new guidance.

² In the interest of determinism and guaranteed behavior we limit ourselves to applications of offline supervised machine learning. Online learning and non-supervised methods introduce additional complications that are easily avoided by adhering to this restriction

Applying software assurance to machine learning: Additional challenges

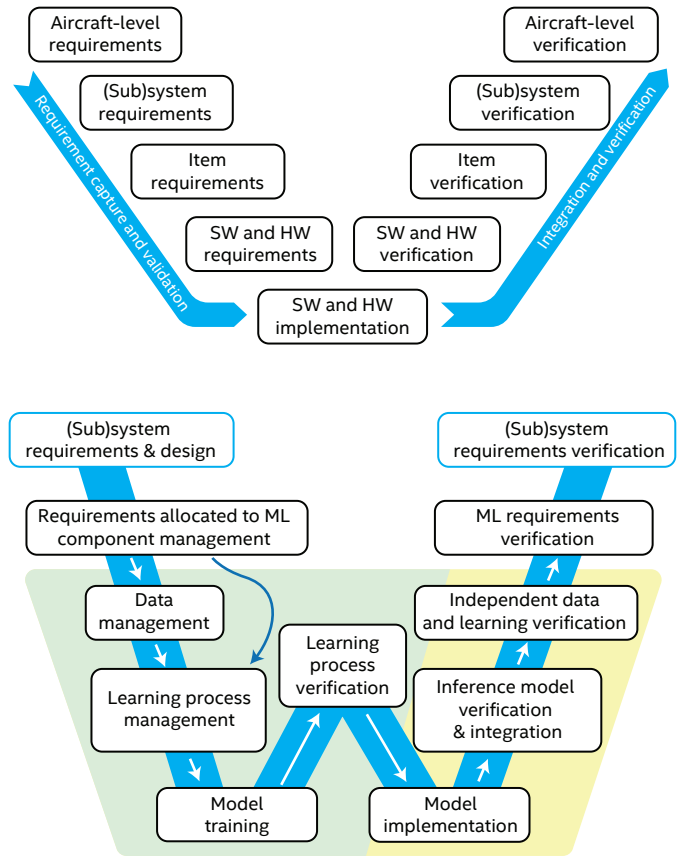
Challenge 1: Traditional software assurance methods are not applicable

DO-178C and related standards strongly emphasize “traceability of requirements” for aircraft systems. Satisfying this exigency involves carefully defining the product’s functional and safety requirements and mapping these requirements to the design and implementation of the software. Tracing these links makes it possible to provide a clear explanation for every line of code in the software and demonstrate how it meets the relevant requirements. Such an explanation is considered an essential aspect of software assurance since it allows regulators to verify that the software meets their standards.

However, the traditional approach to software assurance is not sufficient when it comes to applications based on neural networks. Because machine-learned models are not explicitly defined in lines of code, it is not possible to verify their functionality in the same way. Instead, what is needed is learning assurance, which involves evaluating the quality and reliability of the machine-learned model itself.

Daedalean and the EASA co-published guidelines for learning assurance in the two-part *Concepts of Design Assurance for Neural Networks (CoDANN I, CoDANN II)*. These two papers, along with *DOT/FAA/TC-21/48 Neural Network Based Runway Landing Guidance for General Aviation Autoland* that Daedalean co-published with the FAA, argued that learning assurance can be met by following the W-shaped development cycle, which summarizes the key activities required for the safe use of neural networks (and, more generally, machine learning models) during operation.

The CoDANN reports informed the *First Usable Guidance for Level 1 machine learning applications* proposed by EASA. However, currently, there are no frameworks accepted as Means of Compliance by regulators for certification purposes, and the first certification projects with the FAA or EASA have to use the Issue Paper/Certification Review Item process.



Design and Learning Assurance Processes

Challenge 2: Unprecedented computational performance required from an aerospace-embedded platform

Applications based on AI/ML and processing computer vision are computationally demanding on a level that is unprecedented.

Computationally intensive applications such as AI-enhanced situational awareness—what we call *Situational Intelligence™*—drive aviation’s need for high-performance embedded computing.

Designing electronics for these applications demands high performance in relatively small, rugged embedded form factors to process sensor data volumes and remain compact and relatively low-power.

So, in addition to being:

- a. small in size and weight;

- b. physically robust and tolerant to high vibration and a wide range of temperatures and other environmental conditions;
- c. consuming little electrical energy or, equivalently, not producing excessive heat, and
- d. safety certifiable,

—these applications are based on processing vast amounts of data in a very short time—that is, TOPS are required to process up to a Gigapixel per second.

No suitable commercial and airworthy hardware previously existed to meet these needs. Except in the defense sector, with flexible safety requirements, there was no need for computational systems in commercial avionics to perform at the level Daedalean currently requires.

The Visual Awareness System

Daedalean's Visual Awareness System (VXS)—a component of our Situational Intelligence™ suite—consists of several 12 MP high-resolution cameras and a high-performance, environmentally tested computer containing various hardware components, including a multi-core processor, FPGA accelerators, and a recorder.

The system employs a discrete architecture to shift computationally demanding tasks away from the CPU, allowing the system to operate at high speeds. Table 1 shows the key features of the VXS system.

Key features of the VXS

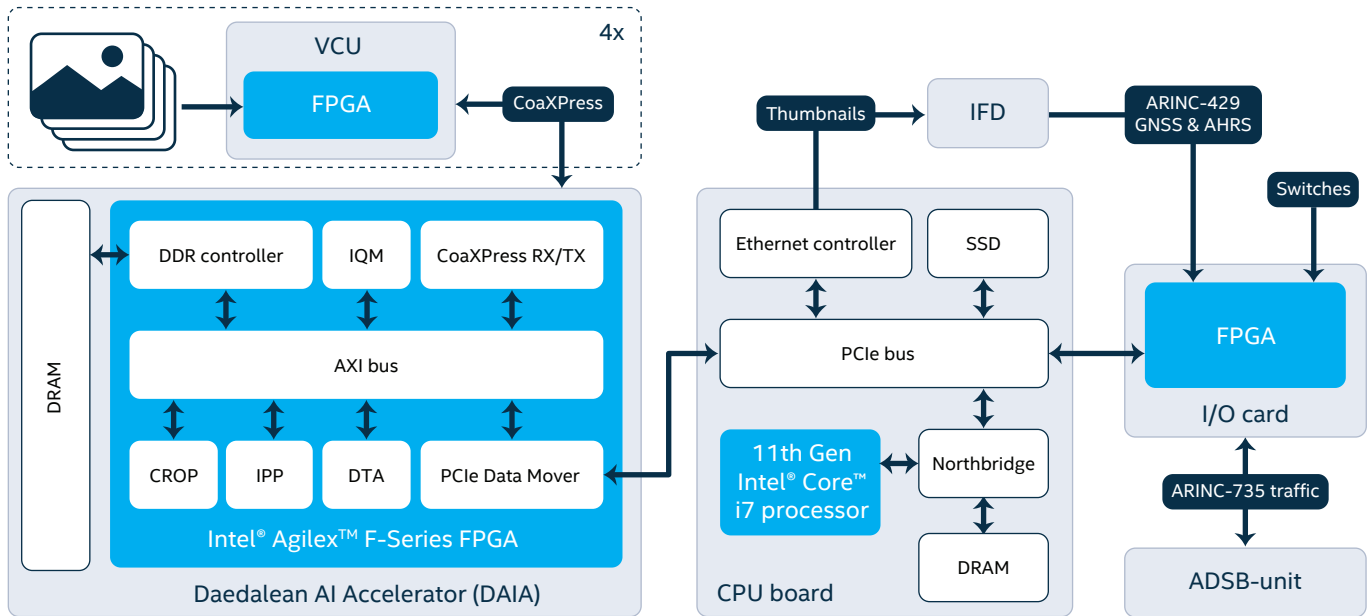
Area	Parameters
CPU	11th Gen Intel® Core™ i7-1186GRE
FPGA	Intel® Agilex™ F-series AGF027 (per 4 cameras, future version) Xilinx® Kintex® Ultrascale+™ KU5P (per camera)
Certification	DO-160G & DO-254, DAL-C with a confirmed path to DAL-B/A
Inputs	From VCU® (Smart Camera)—CoaXPress From IFD / avionics—A429 / A825 or serial DIS—discrete inputs, e.g., weight on wheels
Outputs	To VCU (Smart Camera)—CoaXPress To IFD / avionics—A429 / A825 or serial Ethernet for debugging/maintenance
SWaP	<5 liters, <3 kg, <100 Watts
Performance	Per camera @6 fps, 500 ms latency Intel® Agilex™ F-Series FPGA ³ : 6.6 TOPS/s @ 670 MHz, 122 GOPS/Watt Xilinx® Kintex® Ultrascale+™: 1.2 TOPS/s @ 500 MHz, 58 GOPS/Watt

The VXS uses Daedalean's in-house hypervisor, lightweight virtualization software, to run its application in multiple partitions on the CPU. Daedalean avoided employing a commercial real-time operating system (RTOS), as together with the Intel primitives, the hypervisor can guarantee sufficient determinism.

The application is run in multiple partitions, each containing a single execution thread running on a single physical core and isolated from all other partitions.

³ Estimated performance

The system has the following main components:



Reference Architecture of Daedalean's Situational Intelligence™ Mission Computer based on the 11th Gen Intel® Core™ i7 Processor and Intel® Agilex™ F-Series FPGA. This architecture can process four camera streams, each capturing 6 fps at 6.6 TOPS/s @ 670 MHz, 122 GOPS/Watt

- Each custom-designed Visual Capture Unit (VCU) consists of an image sensor and a MEMS accelerometer/gyroscope.
 - The vcuFPGA controls the sensors and sends the data out through the high-speed CoaXPress link to the “Daedalean AI Accelerator” (DAIA) card (powered by the Intel® Agilex™ F-Series FPGA).
- The DAIA pre-processes the image, splits it into tiles, processes the tiles in the NN, and sends the outputs to the CPU board.
 - The IQM block compiles brightness and contrast on the images as they flow by. Daedalean’s Tensor Accelerator (DTA) performs ML calculations and consumes about 90% of the resources used on the FPGA.
- A data recorder (not shown on the diagram) receives raw, full-resolution images from the camera and stores them for debugging and development purposes.
- The connection to the host, which is the CPU Board, with the 11th Gen Intel® Core™ i7 processor.
- The external controller is used for sending out the output of the system (e.g., traffic information and images about this traffic) to the IFD (integrated flight display).
- The I/O card gathers auxiliary input from other (optional) onboard systems and sends out the actual traffic info to an ADS-B capable traffic advisory system.

The future Intel Agilex-based VXS can handle up to four camera streams on a single chip as opposed to the Xilinx® Kintex® Ultrascale+™ implementation, which can only manage one camera stream per chip. Reducing the four-line cards to one single card provides estimated SWaP improvements of more than 50%. Daedalean is partnering with COTS solutions providers to develop this system, whose performance can be scaled to process additional sensors. It can also be designed to meet open standards, such as Modular Open Systems Approach (MOSA), and Sensor Open Systems Architecture (SOSA™).

To handle the camera stream rate of 12 MP per frame, the VXS® system is based on a proprietary implementation of the public specification called Apache Tensor Virtual Machine—an end-to-end machine learning compiler framework for CPUs, GPUs, and accelerators—which includes a component called [Versatile Tensor Accelerator](#). This component allows for performing ML computations on the camera data, meeting various safety, performance, and design standards, including DO-178C and DO-254.

The VXS is also designed to meet other standards, including DO-365–Minimum Operational Performance Standards (MOPS) for Detect and Avoid Systems, DO-366–MOPS for Air-to-Air Radar for Traffic Surveillance, and DO-387–MOPS for Electro-Optical/Infrared (EO/IR) Sensor Systems for Traffic Surveillance—for technical systems and radars, as well as DO-160–Environmental Conditions and Test Procedures for Airborne Equipment. These standards ensure that the system is fit for purpose and safe to use.

Daedalean Tensor Accelerator

Daedalean Tensor Accelerator (DTA) is an in-house-developed accelerator for deep learning applications based on the Versatile Tensor Accelerator architecture. The DTA was architected and developed by Daedalean and designed for optimal performance in both FPGA and ASIC technologies. It can execute Deep Convolutional Neural Networks of varying topologies and sizes and is flexible enough to run different networks on the same image within a processing cycle. Thanks to architecture-specific optimization, the proprietary FPGA IP currently executes at 550 MHz, close to the theoretical maximum speed rating for the presently used Xilinx family. Emulations of the DTA on the Intel® Agilex™ FPGA estimate that we can achieve 670 Mhz, or at least a 24% speed-up. The scheduling quality of operations on the DTA is done

by the proprietary compiler statically at system build time and can keep it constantly busy to fulfill the requirements of throughput and latency. Processing a single camera stream of 6 frames per second (fps) at a latency of 0.5 seconds requires approximately 1 TOPS, roughly 2x the power provided by the CPU's built-in GPU or of a fully dedicated CPU core.

A framerate of 6 frames per second means that the main scheduling period is approximately 166 milliseconds, with a latency requirement of 500 milliseconds from image capture to outputting the traffic alert.

The system has four cores, but the first product will utilize only 25% of the CPU resources, leaving effectively three complete cores free for future functionality.

The unique advantage of the 11th Gen Intel® Core™ i7 processor: Allowing certification all the way to DAL-A

The introduction of multi-core processor (MCP) architectures has provided performance gains for enterprise general-purpose applications; it has also presented some unique challenges for their use in safety-critical avionics systems. Avionics applications often have specific requirements, including, but not limited to, application isolation and determinism. These are not the primary considerations of semiconductor manufacturers when designing MCPs for the commercial market, which typically pushes for the best average performance. Research undertaken by academia, the avionics industry, and safety certification authorities has found variations between MCP designs. These variations can affect their suitability for use in avionics applications due to the impact of architectural design features on application isolation and determinism.

COTS processors are popular with aerospace suppliers because they offer low-cost alternatives in compliance with military-specified standards, such as MIL-STD-810G, and reduce time to market and production costs. However, demonstrating design assurance of COTS processors can pose severe difficulties since their inner workings are considered proprietary and best known by the silicon vendor. A common design assurance challenge is establishing an application's worst-case performance and deterministic response. To address this challenge, system designers need information about the processor architecture, which they typically cannot access. Without understanding the behavior of multi-core processors, system developers can struggle to manage interference and guarantee mitigation of all potential failure conditions. For example, processor manufacturers may typically withhold details about how a multi-core processor's shared cache works (e.g., when cache lines are evicted), despite the common understanding that cache operation significantly impacts application performance. Consequently, it is a significant advantage for aerospace suppliers to cooperate closely with COTS silicon manufacturers to characterize application performance and behavior rigorously.

The AC/AMC 20-152A, FAA CAST 32A, and the AMC 20-193 were introduced by the EASA and FAA to provide guidance for meeting design assurance of COTS and multicore architectures. Much of the objectives described can only be met through evidence and documentation provided by the component manufacturer. Intel can help aerospace suppliers in their efforts to achieve certifications for flight safety by providing documentation on how the silicon performs under various conditions. Daedalean has developed solutions using modern hardware, such as the 11th Gen Intel® Core™ i7-1186GRE processor and the Intel® Agilex™ line of FPGAs, that enable high-performance tensor computations and can be certified for use up to the highest level of design assurance.

With significant experience in functional safety, Intel has introduced the Intel® Airworthiness Evidence Package for the Intel® Atom® C3708, Intel® Core™ i7-1185GRE and i7-1186GRE, and Intel® Xeon® D-1715TER and D-1746TER processors. The Intel® AEP provides aircraft-embedded manufacturers with processor artifacts to support DO-254 certification up to DAL-A. Intel licenses this data package to help aerospace OEMs achieve airworthiness certification with less development effort and risk than if OEMs attempted to characterize and generate processor artifacts on their own.

In addition to the 11th Gen Intel® Core™ i7, Daedalean is partnering with OEMs and COTS solutions providers to design-in the Intel® Agilex™ F-Series line of FPGAs, which allows implementation of its neural network accelerator at a clock speed of 600 MHz and power consumption of 20W. The Intel® Agilex™ also provides support for CoaXPress I/O and other pixel-level operations.

One requirement mandated by the DO-254 is the analysis and mitigation of single-event effects (SEEs). A SEE safety analysis must be performed as part of the system-level safety assessment for the contribution of the system to catastrophic and/or major failure conditions, i.e., DAL-C and above. This activity is prohibitively expensive to OEMs and, therefore, must be disclosed by the component manufacturer. SEE analysis is typically a standard exercise conducted by chip manufacturers but is rarely disclosed, as its results can infer trade secrets and expose highly sensitive intellectual property. However, Intel provides single-event error reports to OEMs

as part of the AEP, enabling developers to determine SEE error rates at the board, sub-system, and system levels and apply mitigation techniques⁴ that ensure acceptable rates with regard to safety objectives.

For instance, a key benefit of the VXS is that it is designed to handle Single Event Upsets (SEUs)—an example transient SEE—without needing a real-time operating system. This feature is possible because Daedalean's system does not run in a super-critical time loop, and it is possible to verify and correct the FPGA bitstream in the event of an SEU.

Resource Management and AMC 20-193/CAST 32A

Daedalean's whole system runs on Vyper, an in-house developed lightweight hypervisor. For each software component, it determines what and when it is allowed to execute. Vyper features include:


- Minimal partitioning hypervisor running on small code size (~1.5k LOC), making it easier to certify;
- Runs on Intel CPUs and in the Bochs emulator;
- Multiple isolated partitions with each partition statically assigned to one of four physical CPU cores;
- Memory
 - two-level paging (via extended page tables) to provide each Partition with its own isolated address space;
 - memory buffers shared between partitions are possible;
- I/O devices
 - isolation of I/O ports;
 - PCIe device remapping via Intel VT-d;
 - interrupt remapping;
 - virtual PCIe host bridge.

Aviation guidance, namely CAST 32A and the more recent AMC 20-193 (which EASA—but not the FAA—has officially adopted as a replacement to CAST 32A), recommend and in some cases mandate that all available resources are partitioned among applications to ensure predictable execution and prevent competition for access. In the case of a system with four CPU cores, Vyper provides robust partitioning of CPU time, cache levels, system memory, and device access.

Vyper utilizes Intel technologies to partition the cores into non-interfering time slices. We use Intel® Time Coordinated Computing settings in BIOS to make execution more deterministic in general, Intel® Virtualization Technology (Virtual Machine Extensions—VMX) to create virtual machines that run the partitions, Intel® Virtualization for Directed I/O (VT-d) to isolate PCIe devices and assign them to partitions. We also use Intel® Resource Director Technology Framework, particularly Cache Allocation Technology, to partition the shared L3 cache. As such, each partition has its own cache and memory, and the only primitive used is a FIFO for communication, which is set up at compile time. This allows the system to operate without interrupts.

Vyper runs a scheduler on each processor. The scheduler loop is responsible for executing software partitions according to a compile-time-defined schedule. There is no dedicated processor for handling partition hypercalls or managing the overall system state. Instead, the schedule on each processor is responsible for all partitions pinned to that processor. Software partitions may run either application logic or driver code; this distinction is transparent to Vyper.

⁴ Mitigation Techniques can be found in DOT/FAA/TC-15/62 *Single Event Effects Mitigation Techniques Report*



Reference architecture for Daedalean's Vyper hypervisor, which enables multiple applications to run on partitioned resources in the 11th Gen Intel® Core® i7 processor

While partitioning of all resources is not necessary for certification up to DAL-C, Daedalean chooses to conduct partitioning for two reasons: for the potential to upgrade to DAL-B in the future and to simplify the testing process. With full partitioning, all software components can be tested independently without interference.

Daedalean does not require the use of an RTOS in its system because the VXS is designed to be lightweight and self-con-

tained without relying on third-party code. This allows the system to operate without the need for insulation from external sources. That said, OEMs using an RTOS or commercial hypervisor can still run Daedalean software, provided that the RTOS supports the 11th Gen Intel® Core™ i7 processor. OEM co-applications can also run on this architecture, as Daedalean reserves CPU cores for this purpose.

Conclusion

In this paper, we proposed a reference design for addressing future certifiable machine-learned avionics systems requiring high-performance computing at low SWaP. We introduced an example vision-based situational awareness system that leverages neural networks with high-resolution, high-throughput camera inputs. We discussed key challenges and constraints limiting silicon selection and provided a reference architecture based on 11th Gen Intel® Core™ i7 and Intel® Agilex™ F-Series FPGA.

Daedalean's solution combines the performance of Intel® processors and FPGAs with the safety and certifiability of Daedalean's own AI-enhanced situational awareness applications. This combination provides a safe and efficient solution for the unique challenges of applying machine learning in safety-critical aerospace applications.



Notices and disclaimer

Performance varies by use, configuration, and other factors.

No product or component can be absolutely secure.

Your costs and results may vary.

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

Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's Global Human Rights Principles. Intel® products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.