WOLF Advanced Technology

# EYES IN THE SKY, INTELLIGENCE AT THE EDGE: HYBRID GPU–FPGA ARCHITECTURES FOR REAL-TIME AIRBORNE EO/IR MISSIONS

## – WHITEPAPER

Written by: Josephine A. Vitella
Strategic Content Specialist, Creative Marketing

WOLF ADVANCED TECHNOLOGY

## INTRODUCTION

Picture this: an airborne search and rescue aircraft locates a missing hiker using thermal imaging. At the same time, onboard AI distinguishes people from wildlife, stabilizes live video, fuses GPS data, and provides responders with actionable intelligence, all while the aircraft continues its mission. Similar demands exist during border surveillance, maritime patrol, and law enforcement operations, where mission success increasingly depends on transforming massive amounts of sensor data into real-time decisions.

This shift reflects the evolution of modern Electro-Optical/Infrared (EO/IR) systems. Once used primarily for visual observation, today's platforms function as intelligent edge computing systems capable of collecting, processing, and interpreting data from multiple high-resolution sensors simultaneously. Advances in imaging technology have improved image quality and expanded operational capability into darkness and degraded visual environments, while AI has enabled automatic object detection, classification, tracking, and decision support.

These advances, however, come with significant engineering challenges. Higher-resolution sensors, increased frame rates, and sophisticated AI models dramatically increase computational demands, requiring mission computers to process multiple synchronized data streams with minimal latency. At the same time, airborne platforms remain constrained by size, weight, power (SWaP), and environmental requirements, leaving little room for inefficient computing architectures.

Hybrid GPU–FPGA architectures have emerged as an effective solution to these competing demands. By combining deterministic FPGA-based sensor acquisition and preprocessing with the massively parallel AI acceleration of GPUs, these architectures enable each processor to perform the tasks for which it is best suited. The result is a balanced computing platform capable of delivering the real-time performance, reliability, and flexibility required for today's airborne EO/IR missions.

This paper explores how hybrid GPU–FPGA architectures address the growing demands of a irborne EO/IR applications and examines the design considerations that enable high-performance edge computing for public safety missions, including law enforcement, search and rescue, border security, maritime surveillance, and disaster response.

## MISSION CHALLENGES

The computational requirements of airborne EO/IR platforms continue to grow as both sensor capabilities and mission complexity increase. Modern aircraft often carry multiple electro-optical cameras, thermal imagers, laser rangefinders, and, in some cases, radar payloads. Each sensor contributes valuable information, but together they generate substantial volumes of data that must be synchronized, processed, and presented to operators without introducing unacceptable latency.

At the same time, mission expectations have expanded considerably. Rather than simply displaying video feeds, today's mission computers are increasingly expected to execute AI algorithms capable of detecting people, vehicles, vessels, infrastructure damage, wildfire hotspots, or suspicious activity automatically. These AI workloads require significant computational resources while maintaining deterministic timing for sensor acquisition and video processing.

Environmental conditions further complicate airborne operations. Public safety aircraft routinely fly in darkness, smoke, haze, rain, fog, and other degraded visual environments where image quality may be compromised. Image enhancement, sensor fusion, and thermal processing therefore become essential for maintaining situational awareness under challenging conditions.

All of these capabilities must also fit within the SWaP-C limitations imposed by helicopters, fixed-wing aircraft, and unmanned aerial systems (UAS). Computing platforms must deliver increasing levels of performance without exceeding available power budgets or thermal capacity, making efficient hardware acceleration critical to future mission success.

## THE ROLE OF NVIDIA GPUS IN AIRBORNE AI

NVIDIA GPUs provide the parallel-computing foundation required to transform airborne sensor data into actionable intelligence. Unlike conventional CPUs, which execute a comparatively limited number of general-purpose processing threads, GPUs contain thousands of CUDA® cores capable of performing image-processing, signal-processing, and mathematical operations simultaneously. This parallel architecture is particularly effective for airborne EO/IR applications, where multiple high-resolution video streams must be analyzed continuously without delaying mission-critical information.

Dedicated Tensor Cores further accelerate the matrix operations used by deep-learning models. This enables airborne mission computers to perform real-time object detection, classification, scene segmentation, multi-object tracking, and image enhancement directly aboard the aircraft. By executing inference at the edge, systems can identify people, vehicles, vessels, thermal signatures, wildfire hotspots, or damaged infrastructure without first transmitting raw imagery to a ground station or cloud service. Onboard processing therefore reduces communications bandwidth requirements and allows operators to receive results even when network connectivity is limited or unavailable.

The GPU can also support several stages of the mission workload simultaneously. In addition to AI inference, CUDA-based processing can accelerate image stabilization, super-resolution, sensor fusion, coordinate transformations, and the generation of geospatial overlays. Dedicated NVENC and NVDEC engines provide hardware-accelerated video encoding and decoding - NVDEC accelerates the ingest of compressed video streams, while NVENC encodes processed imagery for recording or transmission without placing the full codec workload on the host CPU. NVIDIA's CUDA-X software libraries and high-performance memory architectures further support AI, embedded computing, and data-intensive image-processing applications.

Selecting an NVIDIA GPU for an airborne platform, however, requires more than choosing the device with the highest peak performance. Power limits, cooling capacity, memory bandwidth, environmental requirements, and available integration space all affect the performance that can be sustained during a mission. WOLF's rugged XMC portfolio provides several options for balancing these requirements. For example, the WOLF-3476 and WOLF-3576 provide dedicated GPU acceleration for AI inference, visualization, image enhancement, and video encoding, while the WOLF-3570 combines an NVIDIA RTX™ 2000 Ada GPU with an integrated FGX2 FPGA and GPUDirect RDMA on a single XMC module. This integrated approach can be particularly valuable in compact airborne mission computers where both sensor acquisition and AI processing must operate within constrained space, power, and cooling limits.

Within the overall EO/IR processing chain, the NVIDIA GPUs are best positioned as the system's analytics and visualization engines. They convert synchronized imagery and mission data into detections, tracks, enhanced video, fused situational awareness, and operator-ready displays. FPGA solutions such as the WOLF-3185 FGX2 provide deterministic sensor acquisition, timing, synchronization, and preprocessing, while the GPU supplies the scalable computational performance required for higher-level AI. Rather than competing for the same workloads, these processors perform complementary roles within the sensor-to-intelligence pipeline. This division of responsibilities forms the basis of the hybrid architecture examined in the following section.

## WHY A HYBRID GPU–FPGA ARCHITECTURE?

No single processor architecture is optimized for every stage of the EO/IR processing pipeline. Instead, modern mission computers increasingly rely on heterogeneous computing, assigning different tasks to specialized processors based on their respective strengths.

FPGAs excel at deterministic, low-latency operations that occur immediately after sensor acquisition. Their hardware-configurable architecture allows them to interface directly with numerous video standards while providing precise timing, synchronization, and preprocessing regardless of computational load. Because these functions are implemented directly in programmable logic, they can operate with highly predictable latency—an essential characteristic for mission-critical sensor systems.

GPUs, by contrast, are designed for massively parallel computation. Thousands of CUDA® cores and dedicated Tensor Cores allow modern GPUs to execute complex AI inference workloads, perform image enhancement, accelerate sensor fusion algorithms, and render operator displays simultaneously. These capabilities make GPUs particularly well suited for extracting actionable intelligence from incoming imagery in real time.

The complementary strengths of these processors form the basis of a hybrid architecture.

### HYBRID PROCESSING RESPONSIBILITIES AND DATA PATH

| FPGA Responsibilities | System Data Path | GPU Responsibilities |
|---|---|---|
| Sensor interface and protocol handling | PCIe connectivity between the FPGA and GPU | AI inference, including object detection and classification |
| Frame acquisition, timestamping, and synchronization | FPGA DMA engine and GPU peer-memory access through a supported PCIe topology | Image enhancement and super-resolution |
| Deterministic preprocessing, format conversion, and application-specific image correction | GPUDirect RDMA transfers, where supported by the FPGA, GPU, drivers, and host platform | Multi-object tracking |
| Frame buffering and preparation | Buffer registration, transfer coordination, and synchronization managed by host software | Sensor-fusion algorithms |
| Video routing and stream management | Reduced host-memory copies and CPU involvement after the transfer path is configured | Mission visualization, geospatial overlays, and operator displays<br>Hardware video decoding and encoding through NVDEC and NVENC |

Rather than competing for the same workloads, the FPGA and GPU operate cooperatively through a system-level PCIe data path. The FPGA performs deterministic sensor acquisition and preprocessing, while supported GPUDirect RDMA transfers allow prepared image buffers to move directly into GPU memory without first being staged through host memory. The GPU then performs higher-level AI analytics, image processing, visualization, and video encoding.

## REFERENCE MISSION PIPELINE

A representative airborne EO/IR processing architecture begins with multiple sensor inputs, including electro-optical and thermal cameras, GPS/INS navigation systems, laser rangefinders, and optional radar. These sensors continuously generate large quantities of image and telemetry data that must first be synchronized before higher-level analytics can occur.

The FPGA serves as the initial processing stage, acquiring image frames from the various sensors, correcting and timestamping incoming data, synchronizing multiple video streams, and preparing them for downstream AI processing. Through technologies such as GPUDirect RDMA, image data can then be transferred directly into GPU memory without unnecessary CPU intervention or intermediate memory copies.

Once the imagery reaches the GPU, AI models perform object detection, classification, tracking, image enhancement, sensor fusion, and visualization. The processed information is subsequently displayed to mission operators while simultaneously being recorded or streamed to other systems as required.

By dividing responsibilities across specialized hardware, this pipeline significantly reduces latency while improving overall system throughput and scalability.

### SIMPLIFIED SENSOR-TO-OPERATOR DATA FLOW

The following flow summarizes how sensor imagery and navigation data move from acquisition through deterministic FPGA processing and GPU-accelerated analytics to operator-facing outputs. It is a functional overview; the exact sequence and division of tasks may vary by sensor suite, platform, and mission requirements.

**EO/IR Sensors + GPS/INS + Laser Rangefinder + Optional Radar**
↓
**FPGA: Frame Grab → Synchronization → Image Correction → Timestamping → GPUDirect RDMA**
↓
**GPU: AI Detection → Tracking → Image Enhancement → Sensor Fusion → Mission Display → Recording/Streaming**

## REDUCING LATENCY WITH GPUDIRECT RDMA

As sensor resolutions and frame rates continue to increase, efficient movement of data between processors becomes just as important as the processing itself. In traditional computing architectures, image frames are typically copied multiple times through CPU-managed memory before reaching the GPU for AI inference. Each transfer consumes memory bandwidth, increases processor overhead, and introduces additional latency.

GPUDirect RDMA addresses this bottleneck by allowing FPGA-generated image data to be transferred directly into GPU memory. Eliminating unnecessary memory copies reduces CPU utilization while allowing AI algorithms to begin processing imagery much sooner.

This direct data path becomes particularly valuable when multiple high-resolution EO and thermal cameras are operating simultaneously. Lower latency not only improves operator responsiveness but also enables more sophisticated AI models to execute within the strict timing requirements of airborne missions.

## APPLICATION: AIRBORNE LAW ENFORCEMENT

Airborne law enforcement missions demand rapid situational awareness across highly dynamic operational environments. Aircraft may be tasked with supporting vehicle pursuits, monitoring large public events, conducting tactical overwatch, assisting border patrol operations, or performing maritime surveillance. Operators must continuously interpret multiple video feeds while simultaneously communicating with personnel on the ground.

Within a hybrid GPU–FPGA architecture, the FPGA provides deterministic acquisition and synchronization of video streams from multiple EO and thermal sensors. Once the imagery has been preprocessed and transferred directly to GPU memory, AI algorithms perform real-time detection of people, vehicles, vessels, and other objects of interest. Additional analytics such as scene segmentation, geospatial overlays, intelligent cueing, and target tracking further enhance operator awareness.

Rather than replacing human decision-making, AI serves as a force multiplier by reducing cognitive workload and directing operator attention toward events that require immediate action.

## APPLICATION: SEARCH AND RESCUE

Search and rescue operations frequently involve locating individuals across large geographic areas under challenging environmental conditions. Victims may be partially concealed by vegetation, difficult terrain, darkness, or adverse weather, making manual observation both time-consuming and physically demanding for flight crews.

A hybrid GPU–FPGA architecture improves search efficiency by combining synchronized electro-optical and thermal imagery with advanced AI analytics. The FPGA aligns incoming sensor data while performing deterministic preprocessing to ensure image consistency before transferring frames directly to GPU memory. AI algorithms then identify people, vessels, life rafts, vehicles, thermal signatures, and other objects of interest that may otherwise be overlooked.

By automatically highlighting probable targets and prioritizing areas requiring operator attention, these systems enable crews to search larger regions more effectively while reducing fatigue during prolonged missions.

## APPLICATION: DISASTER RESPONSE

Natural disasters often require emergency responders to rapidly assess evolving situations across large and inaccessible areas. Whether responding to wildfires, floods, hurricanes, or infrastructure failures, decision-makers depend on timely aerial intelligence to prioritize resources and coordinate response efforts.

Hybrid GPU–FPGA architectures enable onboard image analytics that identify wildfire hotspots, map flood extents, inspect damaged infrastructure, and assess changing conditions in near real time. Because this processing occurs directly aboard the aircraft, actionable intelligence can be delivered immediately without relying on high-bandwidth communications or cloud-based infrastructure.

This capability is particularly valuable in disaster environments where communications networks may be degraded or unavailable.

## SYSTEM BENEFITS

The combination of FPGA determinism and GPU acceleration provides several operational advantages over traditional CPU-centric mission computers. Deterministic sensor acquisition ensures that video streams remain synchronized even as AI workloads increase, while GPU acceleration dramatically improves the performance of object detection, tracking, and image enhancement algorithms.

By transferring data directly between the FPGA and GPU through GPUDirect RDMA, CPU utilization is reduced and overall system latency decreases. This architecture also accommodates a wide variety of video interfaces and sensor types, allowing system integrators to support evolving mission requirements without redesigning the underlying computing platform.

Perhaps most importantly, hybrid architectures scale effectively across helicopters, fixed-wing aircraft, unmanned aerial systems, and other edge-computing platforms while remaining compatible with the demanding SWaP-C constraints common to airborne applications.

## REPRESENTATIVE WOLF SOLUTIONS

WOLF's rugged embedded computing portfolio provides several technologies that naturally support hybrid GPU–FPGA mission architectures.

| Product | Primary Role | Best Suited For | Key Highlights |
|---|---|---|---|
| WOLF-3570 | Integrated GPU + FGX2 FPGA | Compact airborne mission computers | NVIDIA RTX™ 2000 Ada GPU, integrated FGX2 FPGA, GPUDirect RDMA, AI acceleration and sensor acquisition on a single XMC module |
| WOLF-3185 FGX2 | FPGA Video Processing | Sensor acquisition and preprocessing | Multi-format video interfaces, deterministic frame capture, synchronization, image correction, and preprocessing |
| WOLF-3476 / WOLF-3576 | GPU Acceleration | AI analytics and visualization | NVIDIA RTX-class CUDA® and Tensor Core processing for AI inference, visualization, image enhancement, and hardware video encoding |

Together, these technologies enable system integrators to build flexible EO/IR mission computers that combine deterministic sensor processing with advanced AI acceleration while meeting the environmental and reliability requirements of airborne deployment.

## FUTURE OPPORTUNITIES

The capabilities of airborne EO/IR systems will continue to evolve as AI models become increasingly sophisticated. Future mission computers are expected to support generative AI for image enhancement, autonomous sensor cueing, collaborative processing across multiple aircraft, and vision-language models capable of assisting operators during complex missions.

As these technologies mature, the volume of sensor data and computational demand will continue to increase. Hybrid GPU–FPGA architectures provide a scalable foundation for supporting these emerging capabilities while preserving the deterministic sensor processing required by mission-critical applications. By separating time-sensitive acquisition tasks from AI-intensive analytics, future systems can continue to integrate increasingly advanced algorithms without compromising reliability or real-time performance.

## PRESERVING OPERATOR TRUST AND SENSOR-DATA INTEGRITY

As generative enhancement, autonomous cueing, and vision-language models enter airborne workflows, their outputs should augment—not obscure—the operator's view of the underlying evidence. Displays should clearly distinguish original sensor imagery from enhanced, fused, inferred, or AI-generated information and should communicate confidence, uncertainty, and relevant system limitations in a form operators can interpret quickly. Mission-critical alerts and recommendations should remain reviewable by a human operator, with straightforward controls to inspect the source imagery, compare enhanced and unaltered views, and override or dismiss automated conclusions.

Maintaining sensor-data integrity also requires an auditable chain of provenance. Systems should preserve original frames and associated metadata; retain timestamps, calibration state, processing steps, and model and software versions; and protect stored and transmitted data against unauthorized alteration. Where imagery may support operational review or evidentiary use, cryptographic integrity checks, access controls, and tamper-evident logs can help demonstrate what the sensor captured and how subsequent processing changed its presentation. Model validation across representative operating conditions, controlled software updates, and ongoing monitoring for degraded or biased performance further ensure that advanced capabilities improve operator effectiveness without weakening confidence in sensor-derived data.

## CONCLUSION

Modern airborne EO/IR mission systems require significantly more than high-performance computing alone. They demand deterministic sensor acquisition, low-latency data movement, advanced AI analytics, and rugged reliability within the SWaP-C constraints of airborne platforms.

Hybrid GPU–FPGA architectures address these requirements by pairing the deterministic processing capabilities of FPGAs with the massively parallel AI acceleration of modern GPUs. Together, these complementary technologies create an efficient sensor-to-intelligence pipeline that improves operator awareness, reduces workload, and enables faster decision-making across law enforcement, search and rescue, border security, maritime surveillance, and disaster response missions.

As airborne public safety operations continue to embrace edge AI and increasingly sophisticated sensor payloads, hybrid GPU–FPGA computing architectures will play a central role in delivering the performance, scalability, and reliability needed for the next generation of EO/IR mission systems.

To learn more about WOLF's cutting-edge solutions,
visit **wolf-at.com** or contact our team at **sales@wolf-at.com**.

WOLF®
ADVANCED TECHNOLOGY

 