

MIL Rugged Scalable SFF AI Compute System

The Cernis small form factor (SFF) NVIDIA Jetson Orin with integrated PCI Express 4.0 switch fabric forms the keystone of an expandable edge AI mission computing system to handle today's highest-end AI applications including sensor fusion, natural language processing, autonomy, situational awareness, and signal intelligence in a SWaP optimized package. While Cernis is a powerful AI mission computer on its own, its real power is unleashed when Cernis is accompanied by up to 24 of the Donati line of NVIDIA Jetson Orin AGX mission computers using an OSS PCIe fabric* to form the lowest latency, highest speed mission compute system. The Cernis/Donati scalable compute system is built from the ground up to tackle multiple performance-hungry AI applications in harsh vehicle environments. With AI inference performance up to 6.6 PetaOPS, data fabric speeds of over 150GB/s and advanced PCIe switching features, the Cernis/Donati scalable AI SFF compute system brings the power of the datacenter to the rugged edge.



The Cernis and Donati products combine in many ways to satisfy a variety of AI processing missions. In 360-degree Situational Awareness (360SA) applications, Donati can be deployed in a backpack or vehicle mounted. Donati-SC acts as a sensor concentrator unit and Donati-MC serves as a GPU accelerated mission computer to serve real time fused video and sensor data to ground forces so nothing is missed outside the vehicle, day or night. With a single Cernis as the centerpiece of the 360SA system, a vehicle mounted system can scale to four

Donati SC (front, rear, left, right) and four Donati MC to provide a full 360SA system for members of the vehicle crew. At this scale the Cernis/Donati system boasts six times greater data throughput and two orders of magnitude lower latency than current generation systems in military vehicle use today. This marked increase in performance allows for sensor fusion, data transport and simultaneous display to rugged crew monitor or tactical helmet display in real time. Since all sensor data is made available to all Donati MCs independently using direct memory access, as if each mission computer owned all of the sensors connected to Cernis, there is no lag or stack overhead caused by traditional networking protocols.



	Cernis	Donati (SC/MC)
Dimensions	14.9" H x 6.51" W x 3.15" D	8.5" H x 4.8" W x 3.15" D
AI Performance	40 TOPS	275 TOPS
GPU	1024-core NVIDIA Ampere architecture GPU with 32 Tensor Cores	2048-core NVIDIA Ampere architecture GPU with 64 Tensor Cores
GPU Max Frequency	625MHz	1.3 GHz
CPU	6-core Arm® Cortex®-A78AE v8.2 64-bit CPU 1.5MB L2 + 4MB L3	12-core Arm® Cortex®-A78AE v8.2 64-bit CPU 3MB L2 + 6MB L3
CPU Max Frequency	1.5 GHz	2.2 GHz
System Memory	8GB 128-bit LPDDR5 68 GB/s	64GB 256-bit LPDDR5 204.8GB/s
Boot Memory	64GB eMMC 5.1	
Network Controllers	2x GbE	
Chassis	Milled Aluminum, Flange Mount	
Weight	19.1 lbs	10.55 lbs
Environment	Temperature: Operating: -40°C to 85°C at 10,000ft altitude Storage: -40°C to 85°C at 10,000ft altitude Humidity: Operating: 0% to 100% condensing relative humidity Storage: 0% to 100% condensing relative humidity Shock: ± 10 a, 11 msec, half-sine pulse, 3 shocks per axis Vibration: 4.5 Grms, 10 to 2000 Hz	
Agency	Designed to conform to the following extended standards: FCC - Verified to comply with Part 15 of the FCC Rules, Class A Canada ICES-003, issue 4, Class A CE Mark (EN55022 Class A, EN55024, EN61000-3-2, EN61000-3-3) CISPR 22, Class A	

*Patent pending

Specifications are subject to change without notice | 8.7.2025