Intelligent Road Infrastructure Solution Blueprint 1.0

Intel ECG – Cities & Critical Infrastructure Division

Contributors:
Prakash Kartha, Anthony Abuta, Vibhu Bithar,
Gustavo Michel, Megha Sarang,
Auguste Muhoza, Yi Qian

Contact us at transportation.us@intel.com

Table of Contents

Executive Summary

The Intelligent Road Infrastructure Solution Blueprint presents a comprehensive framework for developing advanced transportation infrastructure systems leveraging state-of-the-art technology. Utilizing a curated set of silicon products from Intel's extensive portfolio, including Intel® Core™, Intel Atom®, and Intel® Xeon® processors, as well as Intel® Arc™ GPUs, complemented by a meticulously selected software suite, this blueprint addresses critical urban challenges such as traffic congestion, road safety hazards, and operational inefficiencies.

This document is intended for urban planners at federal, state, and local agencies and system integrators. It provides an in-depth analysis including market insights, specific use cases, system architecture, platform dimensioning, performance metrics, and strategic hardware and software partnerships. It offers a detailed strategic plan and a how-to guide for system integrators seeking to modernize transportation infrastructure and enables city planners to make informed decisions in their pursuit of enhanced urban mobility and safety.

Date	Revision	Description
September 2025	2025.2	Initial Release

1. Overview

Intelligent Road Infrastructure systems harness innovative Intel® technologies to advance traffic management, enhance road safety, and optimize urban mobility. These sophisticated systems employ real-time video analytics at critical points such as intersections, highways, and toll plazas to identify and respond to incidents including near misses, red light violations, and wrong-way driving.

The Intelligent Road Infrastructure Solution Blueprint serves as a comprehensive guide for implementing Edge AI solutions that seamlessly integrate with data center and cloud computing. It outlines best practices for leveraging Intel's qualified AI systems, software development kits (SDKs), sample applications, and platforms blueprints. This blueprint not only documents key technologies and system benchmarks but also demonstrates the effective utilization of hardware platforms and software frameworks. By providing an integrated approach that combines real-time event detection and classification at the edge with near-real-time forensic video analysis at traffic command centers, these systems address a multitude of urban challenges.

The blueprint showcases the practical viability and effectiveness of incorporating advanced AI systems and software applications into existing infrastructure. It leverages state-of-the-art AI and video analytics technologies on the latest Intel® Core™, Intel Atom®, and Intel® Xeon® processors and Intel® Arc™ GPUs to demonstrate the transformative potential of Intelligent Road Infrastructure in enhancing urban mobility and citizen safety.

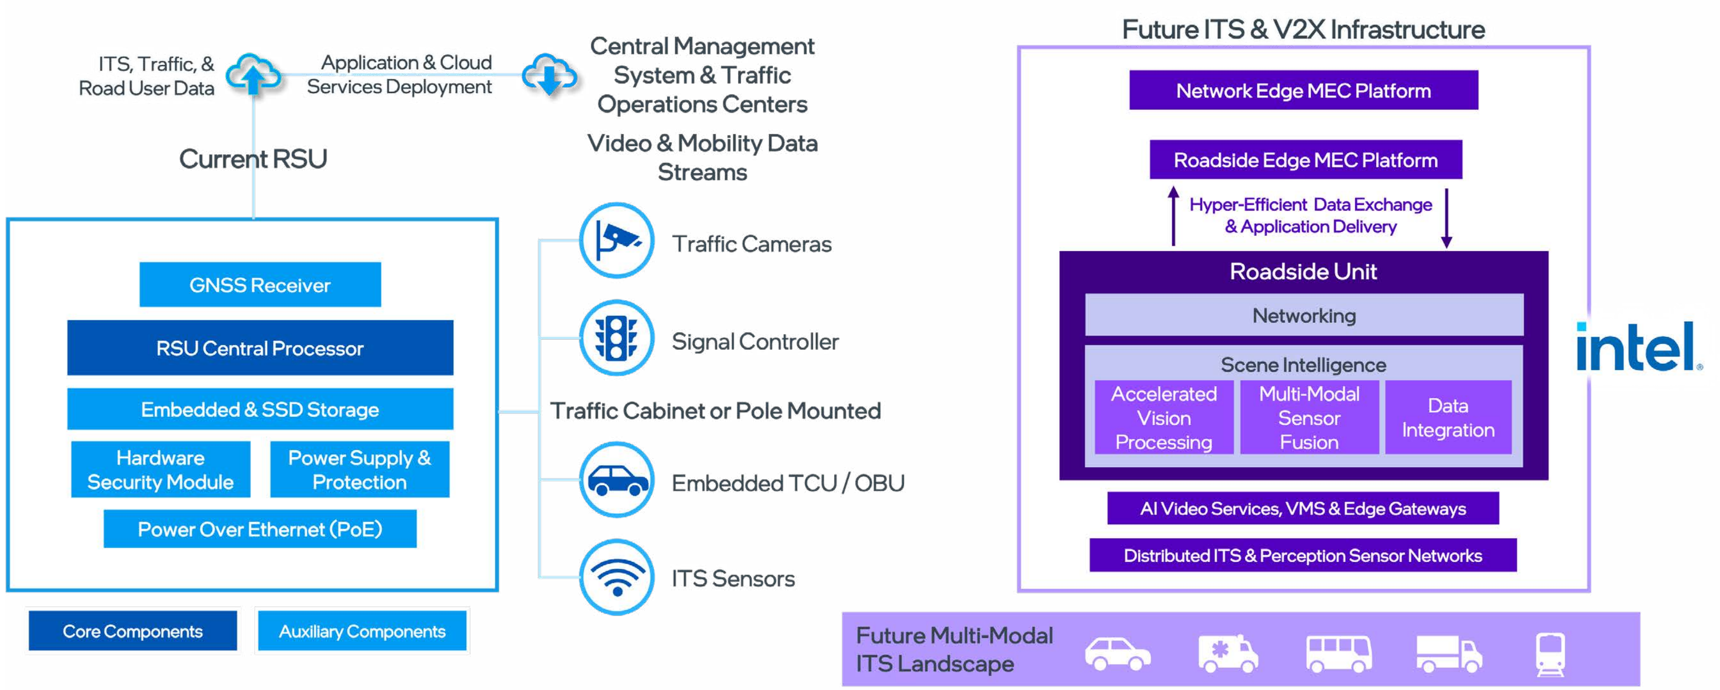

Figure 1. Evolution of the Smart Roadside Unit (RSU) [1]

1.1 Market Analysis

The market potential for Intelligent Road Infrastructure systems is significant, driven by urbanization and the need to improve safety for pedestrians and cyclists while reducing traffic congestion and fatalities. As autonomous vehicles and intelligent traffic systems advance, Smart Roadside Units (RSUs) are crucial for enabling AI-driven edge intelligence, enhancing future mobility.

RSUs, which manage data processing at the far edge, are set to see shipment grow by 31.4% CAGR through 2032, supported by public funding, private investment, and collaborative market efforts focused on improving transportation safety and efficiency [1].

Traffic Management Centers (TMCs) powered by Intel Xeon-class servers to workstation-grade hardware with Graphics Processing Units (GPUs) are projected to grow with the adoption of GenAI/Agentic AI workloads as cities adopt more technology-driven solutions to improve road efficiency, safety, and responsiveness [2].

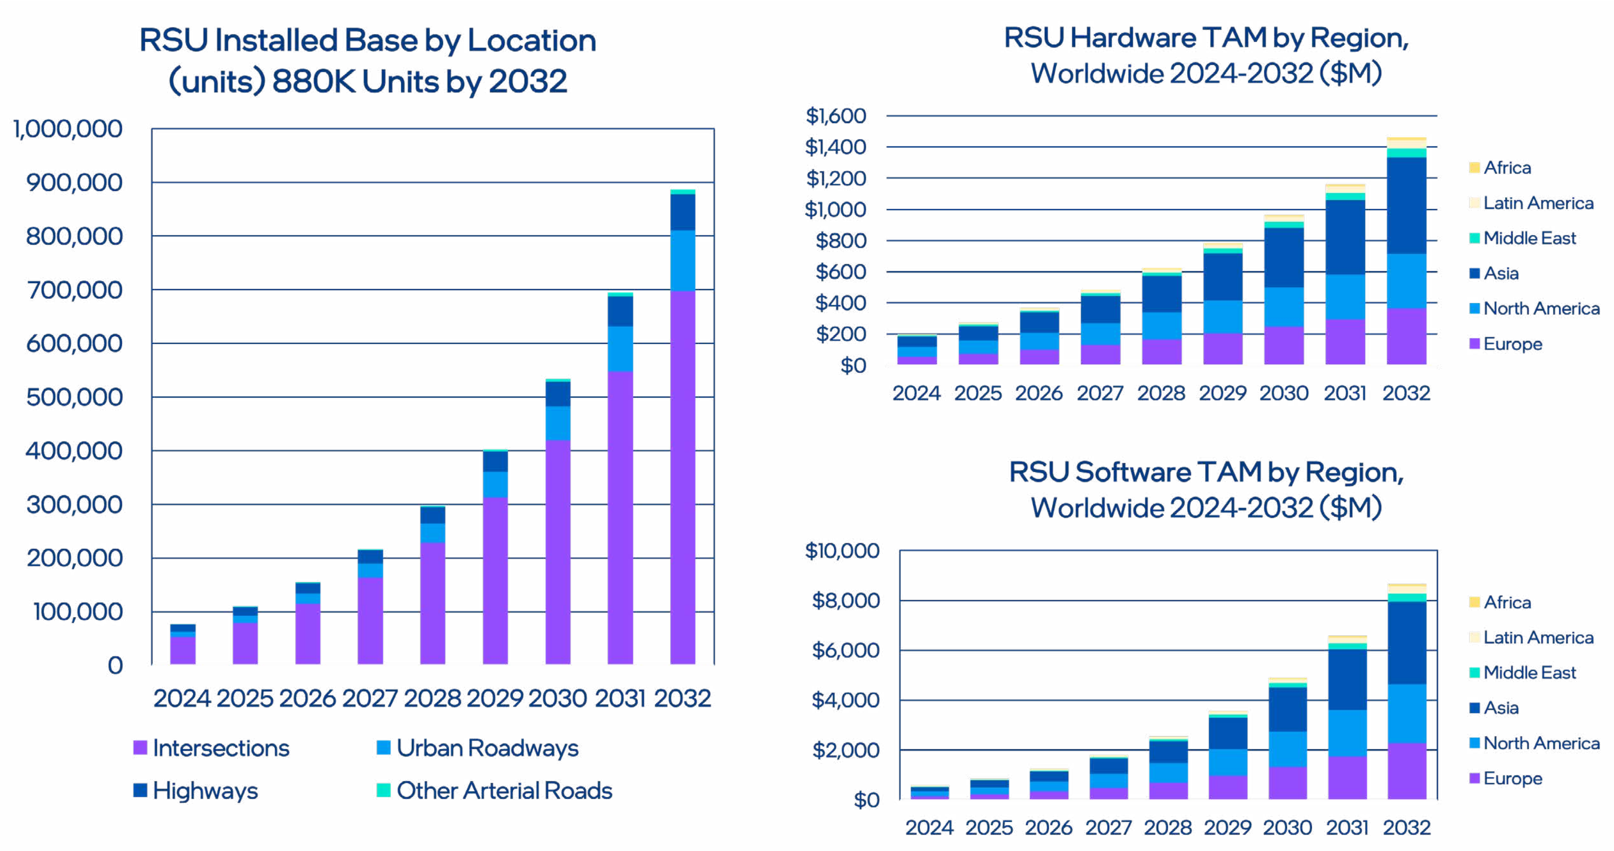

Figure 2. Smart Roadside Unit (RSU) Market Analysis [2]

2. Use Cases

The Intelligent Road Infrastructure System is an advanced solution that combines a Smart Roadside Unit (RSU) with a backend Traffic Forensics System housed in a central control center. This Smart RSU is a robust, high-performance system designed to seamlessly interface with roadside equipment, including intersection and highway cameras, sensors, and tolling systems. It utilizes key technologies such as AI inference, scene intelligence, sensor fusion from multimodal sensors like cameras and radar, generative AI, and cybersecurity, all of which are essential for modernizing infrastructure, enhancing public safety, improving urban mobility, and developing smart roads.

This blueprint supports a variety of use cases, including intelligent traffic management and road condition monitoring. Built to withstand extreme temperatures, the smart RSU offers scalable AI processing capabilities across multiple city intersections and highways, ensuring efficient and reliable operation in diverse environmental conditions. While the listed use cases are among the top applications, they represent just a subset of the potential supported by this technology. With generative AI, the possibilities for additional use cases can be vast.

Traffic Management and Optimization: Smart RSUs gather and analyze real-time traffic data to optimize signal timings and manage lane usage. This dynamic adjustment helps alleviate congestion, improve traffic flow, and reduce travel times in urban areas. Older traffic management systems primarily rely on fixed-timing traffic signals and manual monitoring, using historical traffic data and simple inductive loop detectors to guide traffic flow with limited adaptability to real-time conditions. It is common for inductive loop systems to fail and require constant on-site maintenance, adding to the cost of the solution and causing traffic congestion. Smart RSUs use AI, edge computing, and real-time data from connected vehicles, sensors, and cameras to dynamically optimize traffic flow, reduce congestion, and respond instantly to changing road condition through adaptive signal control and centralized traffic platform. By being able to remotely manage these systems from a control center, cities can increase the uptime and reduce maintenance costs.

Safety and Incident Detection: Smart RSUs enhance road safety by monitoring incidents surrounding vehicle accidents, bicycle

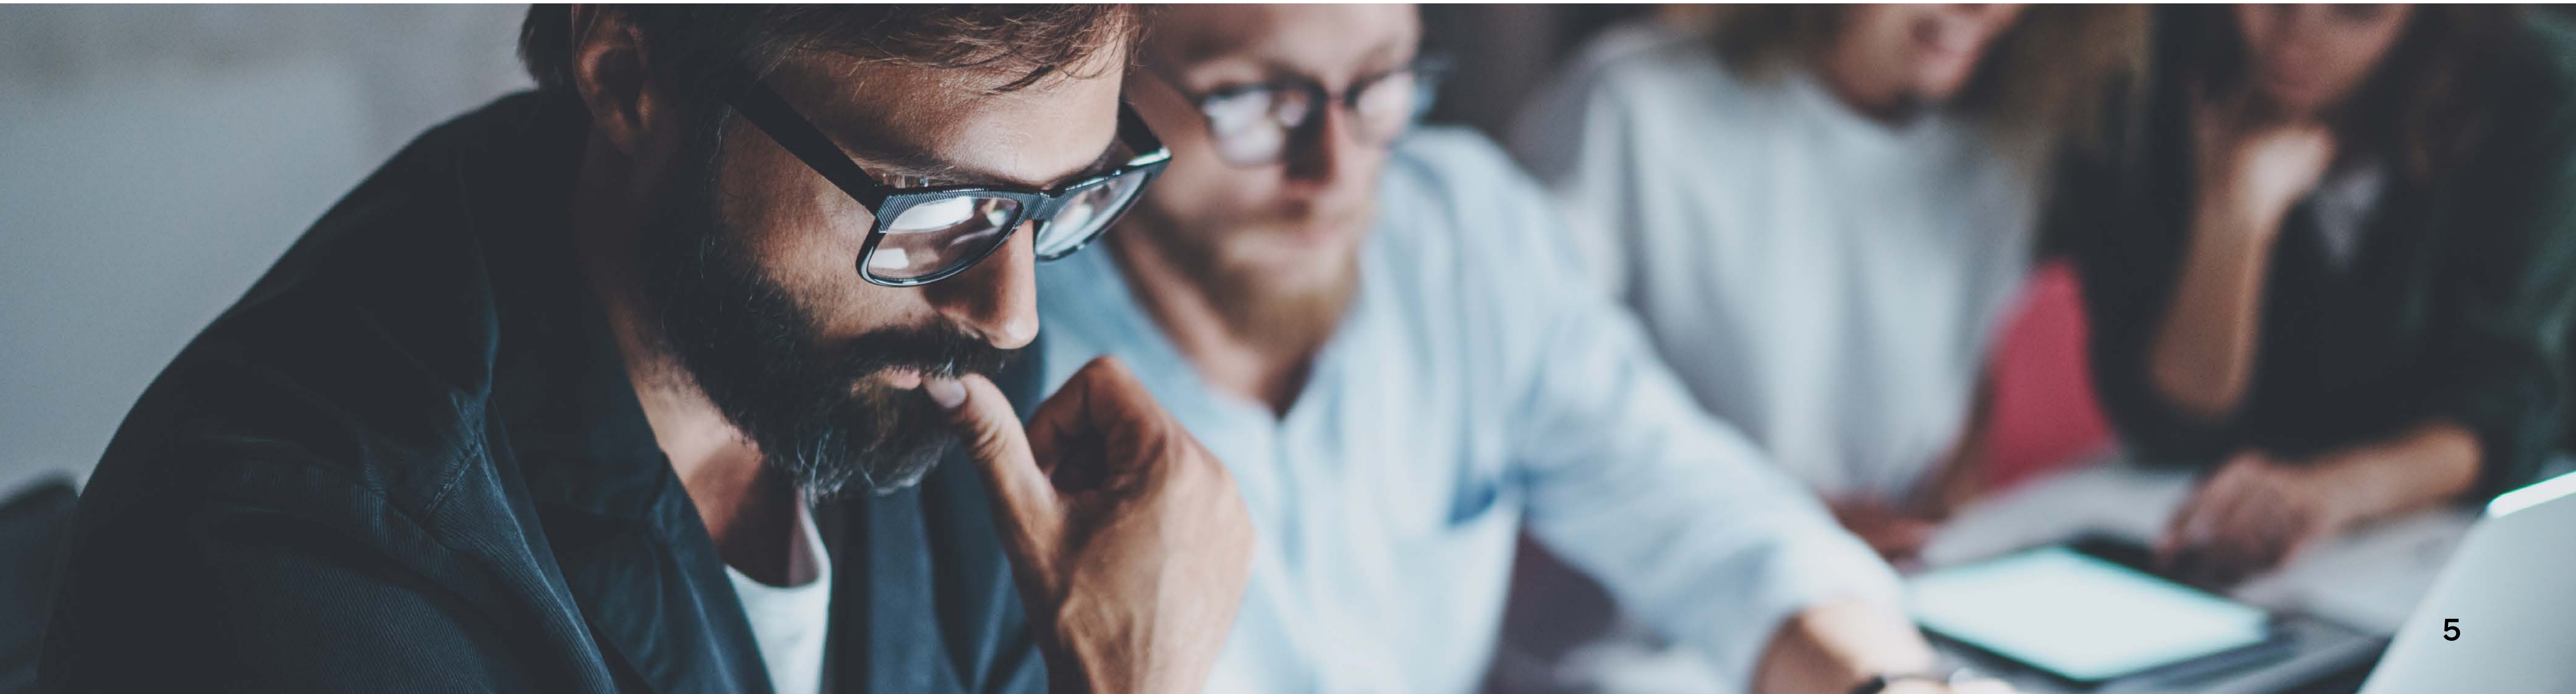

safety, and vulnerable road users and implementing prioritization of emergency vehicles. The data collected at the intersection can be used to analyze such incidents, predict future incidents before they happen, and create early warning signs in real time to alert the travelers and potentially prevent further incidents. These actionable insights can also be provided to traffic management centers and law enforcement agencies, enabling rapid response to the hazardous situation.

Infrastructure Financing and Law Enforcement: Smart RSUs facilitate electronic tolling by using video analytics for license plate recognition, enabling seamless vehicle identification. This can be used to implement dynamic toll pricing, where drivers are charged based on the time of the day and vehicle attributes (size, number of axles, etc.). This reduces congestion at toll booths, enhances traffic flow, and provides a more efficient and convenient tolling experience for drivers. This system can also augment speed detection systems to gather detailed evidence.

Smart Urban Planning: Intelligent Road infrastructure solutions offer transformative capabilities for urban planning, significantly reducing the time required for planning and data collection from months to mere days. By utilizing a recommended hardware and software portfolio, including advanced sensors, AI-driven analytics, and real-time data-processing platforms, urban planners can swiftly gather and analyze comprehensive traffic and environmental data. This setup enables planners to visualize current conditions, predict future trends, and make informed decisions about infrastructure development and resource allocation. The integration of spatial intelligence and IoT connectivity further enhances the ability to model urban environments dynamically, allowing for rapid adjustments and optimizations. As a result, Intelligent Road Infrastructure solutions streamline the urban planning process, facilitating more efficient and effective development strategies that can adapt to the evolving needs of modern cities.

Benefits/Outcomes to Cities

- Reduce travel times, fuel consumption, and emissions
- Improve health and safety with fewer accidents and faster incident response times
- Achieve a rapid return on investment with cost savings and new revenue streams
- Optimize city operations with unified, future-proof systems providing powerful new data

Key Use Cases

- Pedestrian Detection
- Vehicle Detection
- License Plate Recognition
- Street Signal Optimization
- Traffic Congestion
- Digital Signage Management
- Crash Detection
- Traffic Load Optimization
- Emergency Vehicle Prioritaztion
- Send Alerts to Drivers
- Streamline Toll Collection
- Real-Time Analysis for Available Parking
- Road Condition Monitoring
- Bike Lane Detection

Figure 3. Use Case and Outcomes

3. System Architecture

We are entering the era of autonomous driving where the current intersection technology might not be sufficient to support use cases like pedestrian safety, worker safety, and congestion management. The increase in the number of fatalities in road accidents and increased congestion is a testament to that. Although current intersections have cameras, they are primarily deployed for surveillance. Lidar and radar technologies are used but typically to provide redundancies for different, independent systems. Smart RSUs demonstrate how Intel's Edge AI technology can enhance existing infrastructure and provide capabilities like sensor fusion to help with workload consolidation. This is critical and essential for building road infrastructure of the future as this is the foundational piece to gather necessary data/events to build advanced applications like generative AI-based video forensics. The Intelligent Road Infrastructure Solution Blueprint depicted below provides a reference architecture to meet these needs.

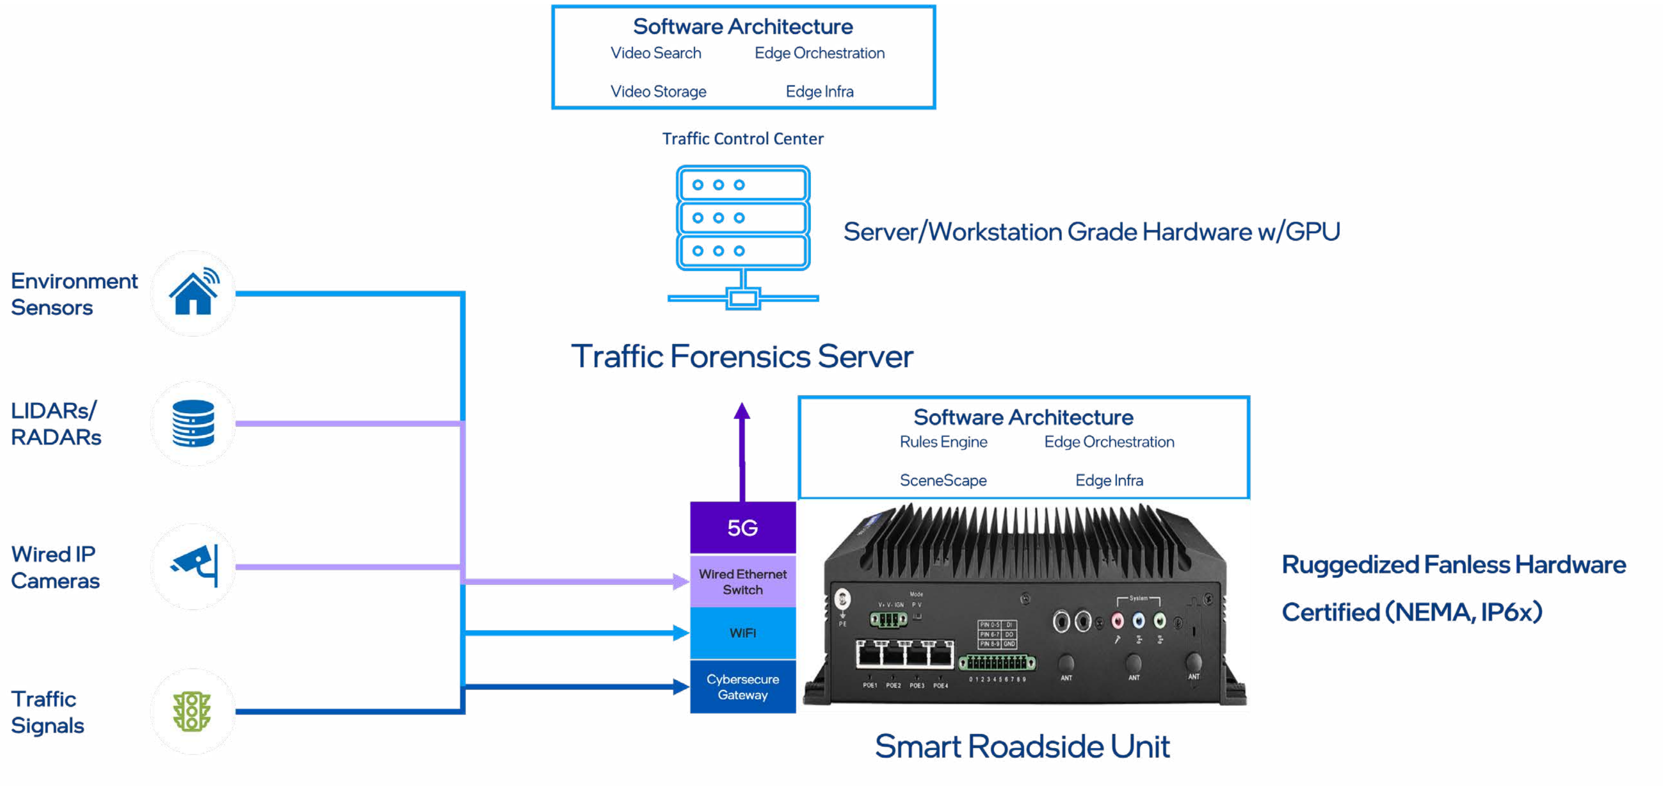

Figure 4. System Architecture

The system topology is composed of two edge nodes:

- **Smart Roadside Unit (RSU):** This is a ruggedized, fanless hardware platform deployed on far edge locations such as intersections, tolling plazas, and highways for real-time sensing of events.
  - Processor Class
    - Intel Core Ultra, Intel Core, or Intel Atom processor.
  - Peripherals:
    - Ethernet switch for IP camera, radar, and lidar sensor inputs.
    - Wi-Fi controller for wireless IoT sensors.
    - Cybersecurity gateway for secure communication with traffic signal controllers.
    - 5G modem for dedicated management and data backhaul to data center or cloud (this is optional if fiber backhaul is available to this system).
  - Operating Temperature Range: -20 °C to 74 °C
  - Software
    - Scene analytics software for vision-based AI and spatial awareness from sensor data (Intel® SceneScape).
    - Rules engine for specifying incident detection criteria.
    - Edge infrastructure software including operating system and various infrastructure services for security and telemetry.
    - Edge orchestration software including remote device and software management.
- **Traffic Forensics Server:** This is a server- or workstation-class hardware platform deployed at the near edge, e.g., a traffic control center, for near-real-time forensics searching and summarization of historical sensor data.
  - Processor Class
    - Intel Xeon or Intel Core Ultra processor.
  - Peripherals
    - Intel Arc GPU for Gen AI video search and summarization.
    - Operating Temperature Range: 18 °C to 27 °C.
  - Software
    - Video Management System (VMS) for storing raw videos.
    - Event monitor/rules engine to monitor events generated by scene analytics and extract video clips from video storage.
    - Vector database for storing embeddings from processed video clips.
    - Natural Language Search and Summarization service based on Vision Language Models (VLM).

Steps in the flow diagram are as follows:

1. Real-Time Streaming Protocol (RTSP) video feeds are sent to the VMS and a scene analytics system like Intel SceneScape.
2. Scene analytics systems analyze the frames and create rich metadata for each object in the frame. Metadata contains details like timestamps, geo coordinates, spatial coordinates, speed, heading, object attributes, etc. Depending on the models used, there can be additional object attributes like license plate, make, model, etc. These details are published over the MQTT bus. Intel SceneScape also generates an event when an object enters a region of interest.
3. The continuous event monitor/rules engine subscribes to event topics and extracts a timestamp from the message and uses that to extract clips from the VMS.
4. Vision language models like CLIP or BLIP are used to extract the vectors from video clips as the video files are received in near-real time.
5. The vectors are then stored in an in-memory vector database like Intel's Visual Data Management System (VDMS).
6. Users can use the natural language queries to search for specific instances.
7. Natural language query is then converted into a vector using CLIP or BLIP.
8. This natural language vector is then matched against a vector in the VDMS to retrieve the associated video. In case the event has not occurred, a user query can still be made into a rule so when a new event happens, a query will be triggered and the user will be alerted.

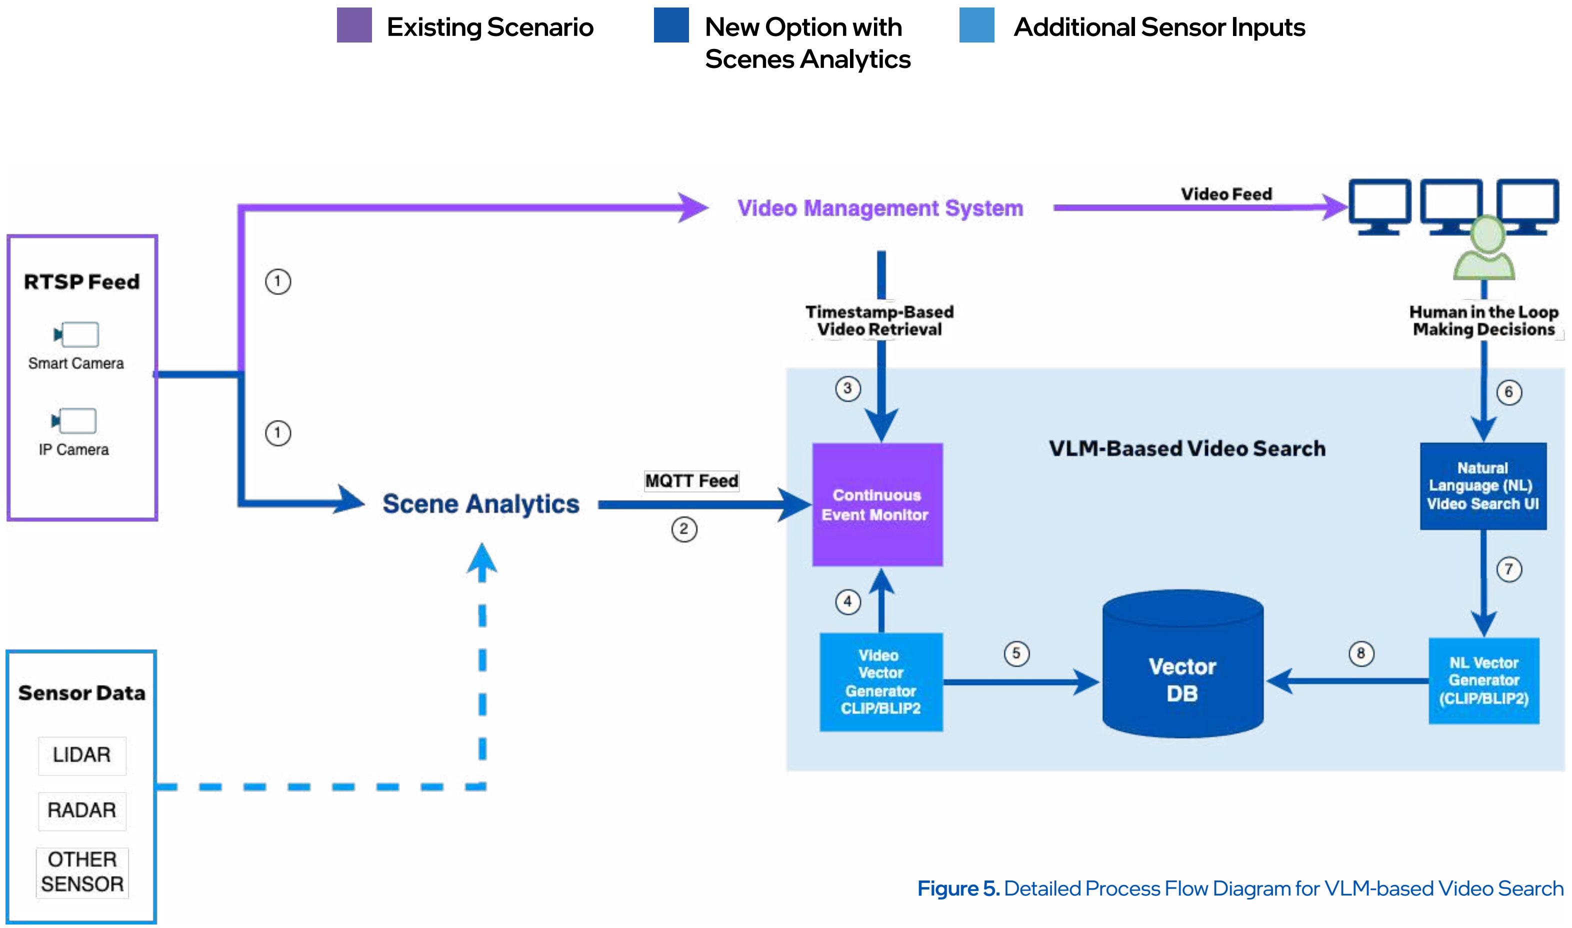

Figure 5. Detailed Process Flow Diagram for VLM-based Video Search

4. Platform Recommendation Guide

4.1 Platform Sizing Methodology

The hardware platform sizing is derived by benchmarking various target silicon platforms with specific end-to-end video/AI pipelines. Each pipeline (see Figure 6 below) is composed of the following components:

- Decode for incoming encoded video streams
- Pre-processing (color conversion, resize, precision conversion)
- AI-based detection of specific objects
- AI-based classification of detected objects
- Optional transcode and display
- Delivery to storage, stream out, or download

Pipelines are created using Intel® Deep Learning Streamer (Intel® DL Streamer) and consistently executed on different platforms and performance measurements captured to derive the platform sizing to match the use case.

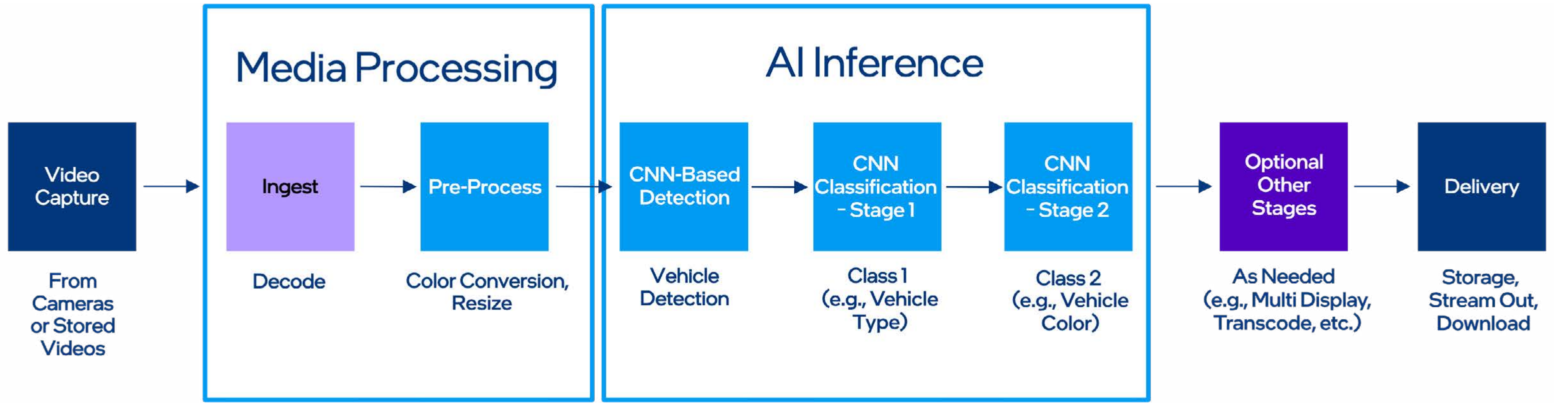

Figure 6. End-to-End Video-Processing and AI Pipeline

4.2 Platform Sizing Guide

Selecting the right hardware for the Intelligent Road Infrastructure Solution hinges on matching the processing power to the demands of each deployment scenario. The blueprint organizes systems into three distinct tiers — Entry, Mainstream, and Performance — each tailored to handle specific workloads, from lightweight monitoring to heavy-duty AI-driven analytics. These tiers account for variables like types of use cases, camera counts, frames per second (fps), end-to-end AI inference, video processing pipeline, and workload requirements.

Figures 7 and 8 below outline each tier's specifications, providing a clear guide for engineers and planners to align hardware with operational needs. Contact us for specific business application needs.

For Smart Roadside Unit Edge Platform:

Entry	Mainstream	Performance
Use Case: Single-gantry tolling, low AI density Smart Intersection applications	**Use Case:** <8 lane multi-gantry tolling, medium-AI-density Smart Intersection applications	**Use Case:** 16+ lane multi-gantry tolling, heavy-AI-density Smart Intersection applications
Video Requirements: 1-4 cameras, ~12 fps, low performance, low accuracy	**Video Requirements:** 4-8 cameras, ~20-30 fps, low-to-medium performance, medium accuracy	**Video Requirements:** 8+ cameras, ~20-30 fps, mid-high performance, high accuracy
Workload: Detect: detect cars and pedestrians Detect-Classify: detect-read license plate	**Workload:** Detect-Track-Detect-Classify: detect and track the vehicle, detect and read the license plate	**Workload:** Detect-Track-Classify-Detect-Classify: Detect and track the vehicle, classify its make/model, detect and read the license plate
Platforms Used: Intel Atom® X-Series processor family	**Platforms Used:** Intel® Core™ processor (14th Gen)	**Platforms Used:** Intel® Core™ Ultra processor (Series 2)

Figure 7. Platform Sizing Guide for RSU

For Video Forensics Edge Server:

Entry	Mainstream	Performance
Use Case: Video search and summarization from low-event density	**Use Case:** Video search and summarization from medium-event density	**Use Case:** Video search and summarization from high-event density
Video Requirements: Process 40 clips per min (e.g., single 4-camera intersection with 10 events per minute per camera)	**Video Requirements:** Process 4000 clips per min (e.g., ten 4-camera intersections with 100 events per minute per camera)	**Video Requirements:** Process 400,000 clips per min (e.g., 100 4-camera intersections with 1000 events per minute per camera)
Workload: 150M parameter VLM for simple video search by text: CLIP (OpenAI)	**Workload:** 3B parameters VLM for advanced video search by text: PaliGemma 3B (Google)	**Workload:** 12B parameter VLM for more advanced video search by text: Pixtral-12B (Mistral)
Platforms Used: Intel® Xeon® Scalable processors + 2 GPUs	**Platforms Used:** Intel® Xeon® Scalable processors + 4 GPUs	**Platforms Used:** Intel® Xeon® Scalable processors + 6 GPUs

Figure 8. Platform Sizing Guide for Video Forensics Server Projected

4.3 Comparative Benchmarks of Intel® Processors

Smart RSU performance is critical for transportation infrastructure. This benchmarking evaluates Intel processors—Core Ultra, Core, and Atom—on ruggedized platforms at far-edge locations (intersections, toll plazas, highways).

Using Intel® Deep Learning Streamer pipelines, the test focuses on License Plate Recognition (LPR) via decodebin3, YOLO-11n-INT8 for detection, and ch_PP OCRv4-FP32 for classification on a headless display. Each stream uses @ 25±2 frames per second (FPS).

Results:

- Intel Atom® x7835RE: 4 streams
- Intel® Core™ i5-1350P: 8 streams
- Intel® Core™ Ultra 9 285H: 18 streams

These results define the Total FPS benchmark for Entry, Mainstream, and Performance platforms in transportation applications.

Please reach out for a copy of the detailed Intel Processor Comparative Benchmarks.

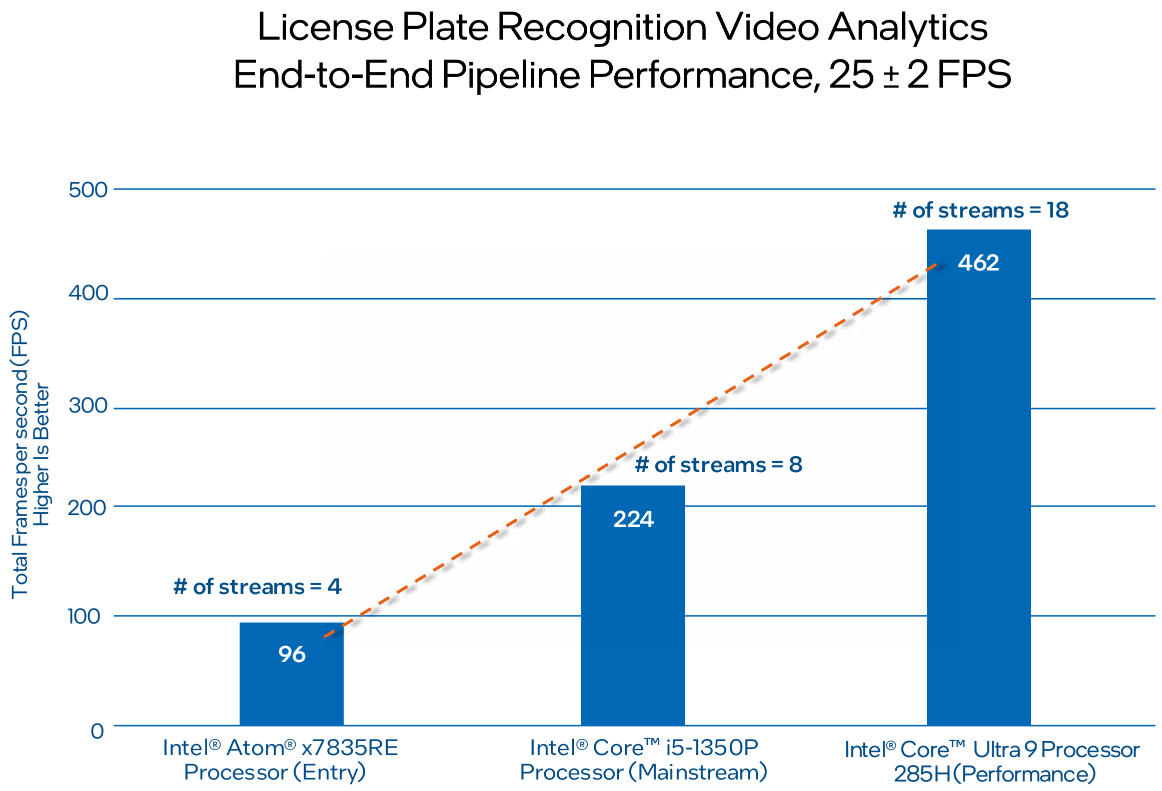

Performance varies by use, configuration and other factors. See Appendix for configuration details.

Figure 9. Intel® Processors Total Frames-per-second Comparison for License Plate Recognition Entry, Mainstream and Performance Transportation Platforms

4.4 Competitive Benchmarks and Analysis

The processor benchmarks reveal Intel's superior performance in end-to-end video processing and AI inference pipeline, offering up to **double** the Edge AI performance of competitors at **half** the price. The benchmarks include HEVC-1080 media decoding at 30 fps and AI workloads with YOLOv5s, ResNet-50, and MobileNet-v2 models, all using INT8 precision and achieving 5 inferences per second. Intel Core Ultra processors provide **significant** cost savings compared to the competitor board, highlighting Intel's competitive edge in end-to-end high-performance, cost-effective video processing, and AI solutions at the edge.

For further information on competitive analysis, please reach out to us for the benchmarking results.

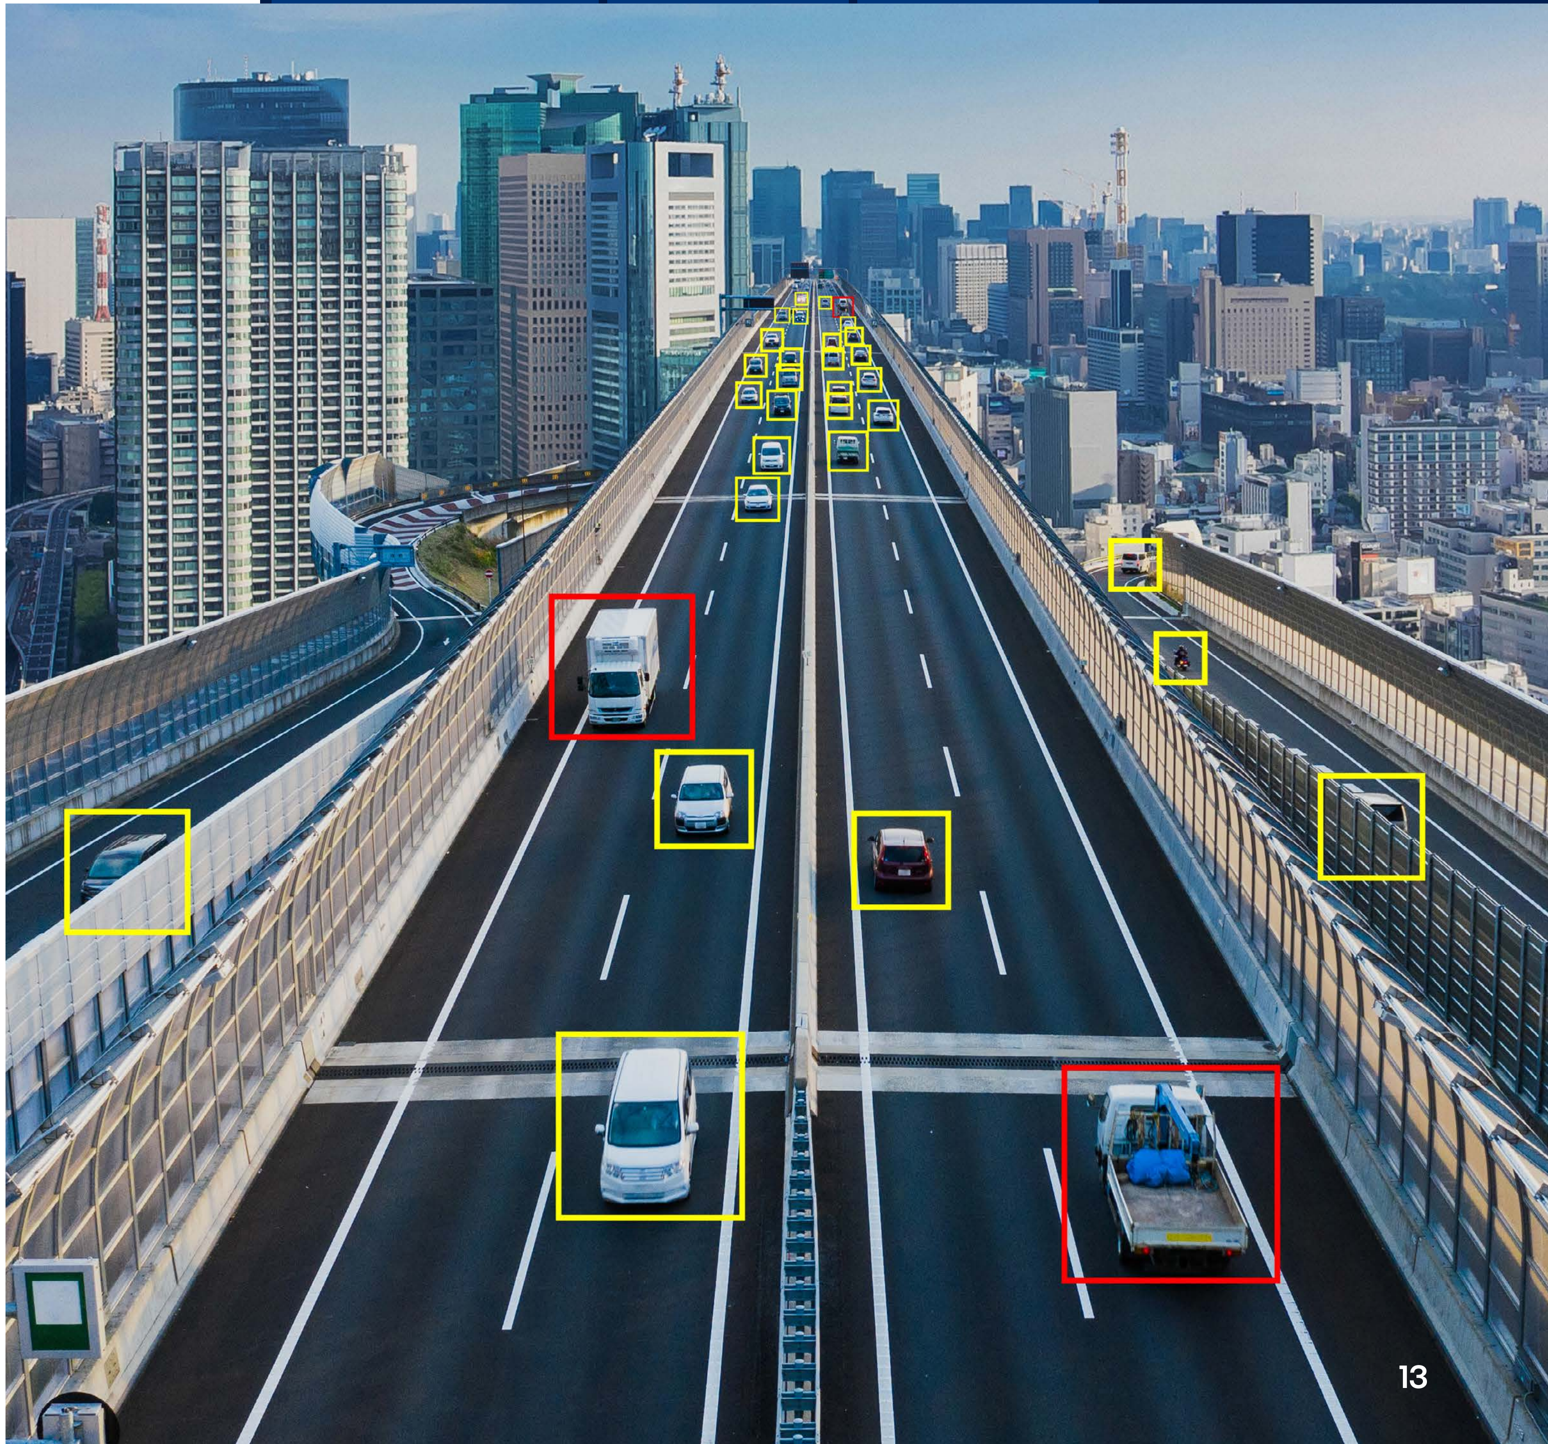

4.5 Certification Requirements Guide

Smart RSUs operate in unforgiving environments — blazing heat, torrential rain, or corrosive salt air — so compliance with industry standards is nonnegotiable. These units must endure years of exposure while maintaining uptime, meeting safety regulations, and ensuring interoperability with other systems. The table below outlines these standards, specifying their purpose, application, and testing criteria to guide deployment teams in certifying RSUs for long-time reliability.

Category	Certification	Purpose
Systems Safety and Functional Certifications	IP66, IP67	Essential for Smart RSUs as they ensure protection against dust and water ingress, allowing the units to operate reliably in harsh outdoor conditions, including heavy rain and temporary submersion.
	IP40	To ensure protection against solid objects larger than 1 mm, ensuring basic protection from dust and small particles in harsh environments.
	NEMA (National Electrical Manufacturers Association)	Ensures Smart RSUs meet standards for safety, reliability, and environmental resilience, allowing them to withstand harsh conditions like dust, water, and extreme temperatures. This certification is essential for RSUs to operate reliably in challenging outdoor environments.

Table 1: Smart RSU Certification Guide

4.6 AI Cybersecurity Technologies

This section outlines Intel technologies that enable AI cybersecurity within the platform roadmap for the Intelligent Road Infrastructure solutions. It demonstrates cyber threats analysis for these solutions and the ways these threats can be mitigated in distributed edge deployments by leveraging Intel technologies.

4.6.1 Threat Analysis

The Intelligent Traffic Management System (ITMS) represents a digital infrastructure that leverages edge computing and AI inferencing to optimize traffic flow and bolster public safety [3]. Typically, it has a distributed edge deployment composed of multiple interconnected elements, as illustrated in Figure 9. This distributed edge deployment with data and AI workloads across multiple edge platforms introduces AI cybersecurity vulnerabilities, underscoring the need for robust AI cybersecurity strategies for Edge AI applications to ensure its integrity and operational reliability. Examples of threats include:

- The integrity of data sources, including camera and sensor feeds, can be compromised through unauthorized alterations [4].
- Edge systems, such as the Smart Roadside Unit or Traffic Forensics Server, are vulnerable to denial-of-service or ransomware attacks, which can disrupt the reception of traffic data feeds.
- Pedestrian and road users' information is at risk of unauthorized modifications for malicious purposes, as well as privacy compromises.
- Actionable insights derived from data analytics can be subject to tampering, affecting decision-making processes.
- The AI model operating within the edge device, responsible for processing video feeds, may be susceptible to compromise, potentially leading to inaccurate analytics, which can lead to undesirable actions.

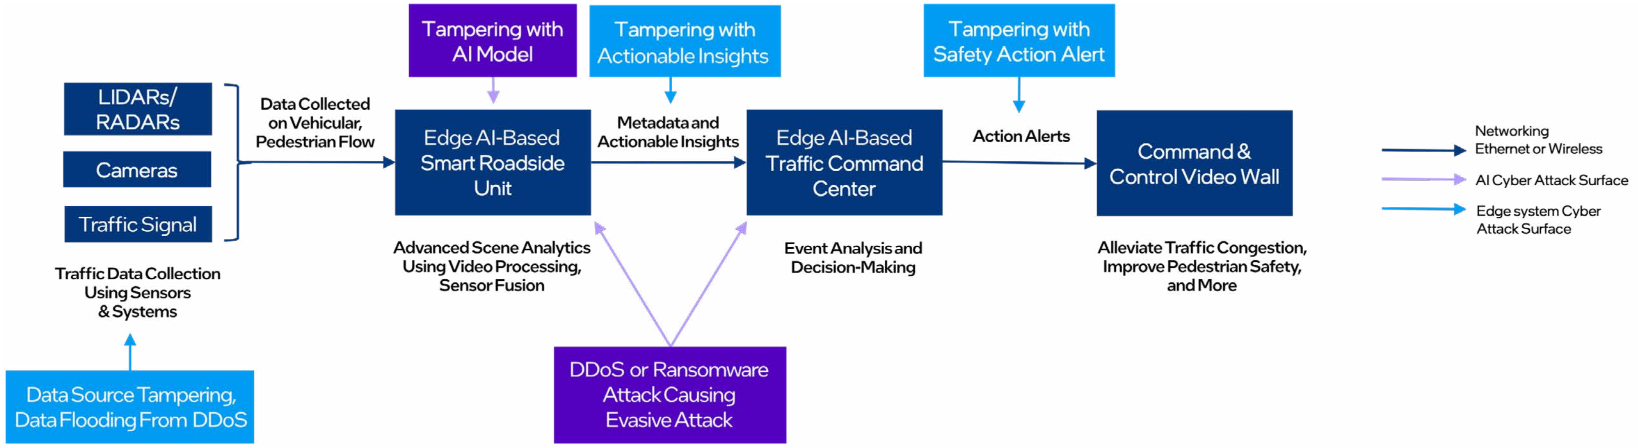

Figure 10. Intelligent Traffic Management System Threat Analysis [4][5]

4.6.2 Intel® Cybersecurity Technologies Across the Road Infrastructure Platform Guide

Based on the Platform Sizing Guide (Section 4.2) for this blueprint, this section outlines Intel technologies that enable AI cybersecurity within this platform roadmap.

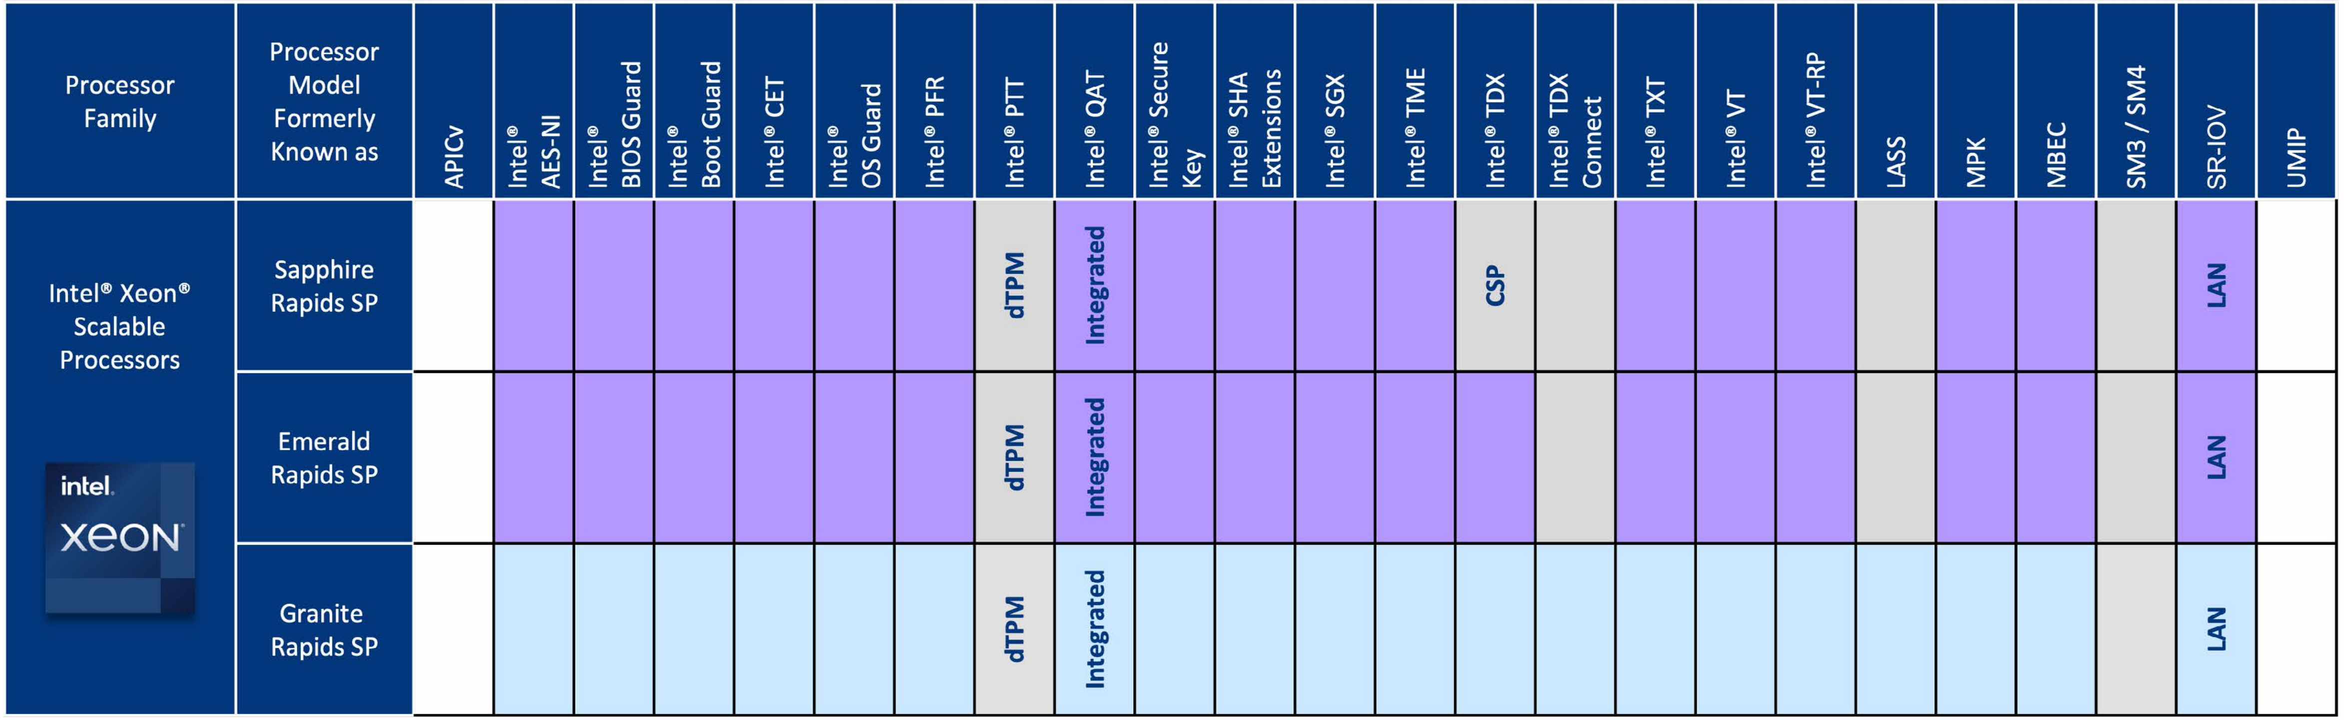

Processor Family	Processor Model Formerly Known as	APICv	Intel® AES-NI	Intel® BIOS Guard	Intel® Boot Guard	Intel® CET	Intel® OS Guard	Intel® PFR	Intel® PTT	Intel® QAT	Intel® Secure Key	Intel® SHA Extensions	Intel® SGX	Intel® TME	Intel® TDX	Intel® TDX Connect	Intel® TXT	Intel® VT	Intel® VT-RP	LASS	MPK	MBEC	SM3 / SM4	SR-IOV	UMIP
Intel® Xeon® Scalable Processors	Sapphire Rapids SP		🟪	🟪	🟪	🟪	🟪	🟪	▨ dTPM	🟪 Integrated	🟪	🟪	🟪	🟪	▨ CSP	▨	🟪	🟪	🟪	▨	🟪	🟪	▨	🟪 LAN	
	Emerald Rapids SP		🟪	🟪	🟪	🟪	🟪	🟪	▨ dTPM	🟪 Integrated	🟪	🟪	🟪	🟪	🟪	▨	🟪	🟪	🟪	▨	🟪	🟪	▨	🟪 LAN	
	Granite Rapids SP		🟦	🟦	🟦	🟦	🟦	🟦	▨ dTPM	🟦 Integrated	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	🟦	▨	🟦 LAN	

Figure 11. Security Features Across Intel® Xeon® SP Platforms in Road Infrastructure Roadmap

🟪 Plan of Record (POR)
🟦 Planned
▨ Unavailable

dTPM: For Intel Xeon-based platforms, there is no support for integrated TPM; use discrete TPM instead.
CSP: In Sapphire Rapids SP, Intel® Trust Domain Extensions (TDX) is available only via cloud service provider offerings, not supported for edge deployments.
LAN: SR-IOV support for select Intel® Ethernet products.

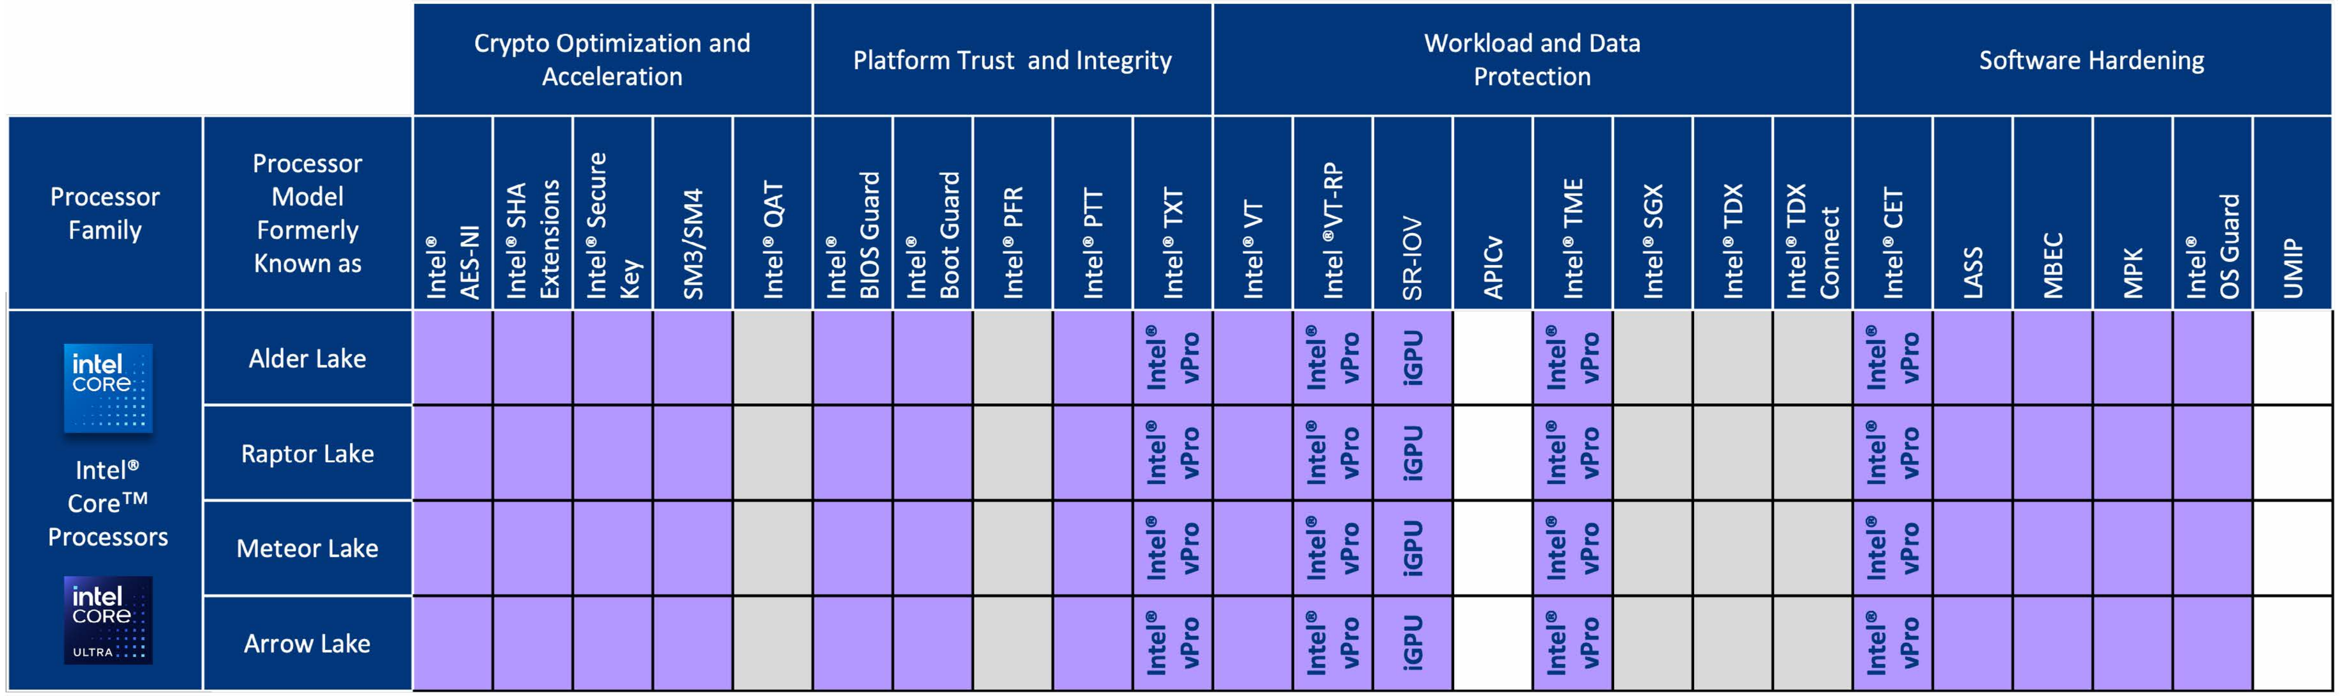

		Crypto Optimization and Acceleration					Platform Trust and Integrity					Workload and Data Protection								Software Hardening					
Processor Family	Processor Model Formerly Known as	Intel® AES-NI	Intel® SHA Extensions	Intel® Secure Key	SM3/SM4	Intel® QAT	Intel® BIOS Guard	Intel® Boot Guard	Intel® PFR	Intel® PTT	Intel® TXT	Intel® VT	Intel ®VT-RP	SR-IOV	APICv	Intel® TME	Intel® SGX	Intel® TDX	Intel® TDX Connect	Intel® CET	LASS	MBEC	MPK	Intel® OS Guard	UMIP
Intel® Core™ Processors	Alder Lake	🟪	🟪	🟪	🟪	▨	🟪	🟪	▨	🟪	Intel® vPro	🟪	Intel® vPro	iGPU		Intel® vPro	▨	▨	▨	Intel® vPro	🟪	🟪	🟪	🟪	
	Raptor Lake	🟪	🟪	🟪	🟪	▨	🟪	🟪	▨	🟪	Intel® vPro	🟪	Intel® vPro	iGPU		Intel® vPro	▨	▨	▨	Intel® vPro	🟪	🟪	🟪	🟪	
	Meteor Lake	🟪	🟪	🟪	🟪	▨	🟪	🟪	▨	🟪	Intel® vPro	🟪	Intel® vPro	iGPU		Intel® vPro	▨	▨	▨	Intel® vPro	🟪	🟪	🟪	🟪	
	Arrow Lake	🟪	🟪	🟪	🟪	▨	🟪	🟪	▨	🟪	Intel® vPro	🟪	Intel® vPro	iGPU		Intel® vPro	▨	▨	▨	Intel® vPro	🟪	🟪	🟪	🟪	

Figure 12. Security Features Across Intel® Core™ and Intel® Core™ Ultra Platforms in Road Infrastructure Roadmap

🟪 Plan of Record (POR)
🟦 Planned
▨ Unavailable

GFX: SR-IOV support for integrated GPUs.

Processor Family	Processor Model Formerly Known as	APICv	Intel® AES-NI	Intel® BIOS Guard	Intel® Boot Guard	Intel® CET	Intel® OS Guard	Intel® PFR	Intel® PTT	Intel® QAT	Intel® Secure Key	Intel® SHA Extensions	Intel® SGX	Intel® TME	Intel® TDX	Intel® TXT	Intel® VT	Intel® VT-RP	LASS	MPK	MBEC	SM3 / SM4	SR-IOV	UMIP
Intel Atom® Processors	Elkhart Lake		🟪	🟪	🟪		🟪	▨	🟪 BOTH	▨	🟪	🟪	▨	▨	▨	▨	🟪		▨	▨	🟪	🟪	▨	
	Amston Lake		🟪	🟪	🟪		🟪	▨	🟪 BOTH	▨	🟪	🟪	▨	▨	▨	▨	🟪		🟪	🟪	🟪	🟪	🟪 GFX	

Figure 13. Security Features Across Intel Atom® Platforms in Road Infrastructure Roadmap

- 🟪 Plan of Record (POR)
- 🟦 Planned
- ▨ Unavailable

GFX: SR-IOV support for integrated GPUs. BOTH: Integrated TPM and discrete TPM supported.

4.6.3 Cyber Threats Mitigation Leveraging Intel® Cybersecurity Technologies

This section revisits the threats identified in the ITMS use case threat analysis (Section 4.6.1) and examines the potential AI cybersecurity attacks each threat may lead to. Further, it highlights how these threats intersect with the use case outlined in Section 4.6.1. Finally, we describe mitigation mechanisms and showcase Intel security technologies that can help prevent the listed attacks.

There are five categories identified for AI cybersecurity threats: data tampering and poisoning (for data both at rest and in transit), attacks on AI models, distributed denial of service (DDoS) attacks, privacy and data breaches, and vulnerabilities in multi-tenant Edge AI workloads within shared edge environments.

The table below outlines the five categories of AI cybersecurity threats, provides a breakdown of the specific attacks we identified within each category, and maps them to the use case (described in Section 4.6.1). It also presents mitigation strategies and the corresponding Intel security technologies.

AI Cybersecurity Threat/Attack Types		Mitigation/Prevention	Intel® Security Technologies
Data Tampering & Poisoning (Data in Transit/at Rest)	Input Data for AI Algorithms Metadata and Insights Action Alerts Data	Encryption	Intel® AES-NI Intel® Secure Key Intel® PTT Intel® QAT
Attacks on AI Models	Modifying Input Data to AI Models (Evasion)	Memory Encryption	Intel® TME
	Altering AI models	AI Model Integrity Protection through Hashing	Intel® SHA MPK
	AI Control Flow Hijacking	AI Runtime Integrity Protection through Memory Corruption Resilience	Intel® CET MPK
DDoS	Overconsume Resources (CPU/GPU/NPU/Memory)	Hypervisor Protection	Intel ® TDX Intel® VT & Intel® VT-d MBEC
	Overwhelm AI algorithms	Protected Execution	Intel® TXT SR-IOV
	Data Flood Mimicking Collected Data	Data Integrity Verification	Intel® TXT Intel® PTT
Vulnerabilities in Multi-Tenant Edge AI Workloads within Shared Edge Environments	AI algorithm Theft Data Privacy Compromise Resources Contention	Workload Isolation Resources Isolation Memory Access Protection Secure and Isolated I/O Access	SR-IOV for Network Cards & GPUs Intel® TDX Intel® VT & Intel® VT-d Intel® SGX Intel® TME MPK MBEC

Table 2: Cyber Threat Types and Mitigation Using Intel® Cybersecurity Technologies

5. Software Development Kits & Tools

This section presents a collection of software development tools and reference software from Metro AI Suite, a part of Intel's Edge AI Portfolio, intended to support various stages of the design and development process for Intelligent Road Infrastructure solutions. These software packages are offered royalty-free and are modular, allowing developers to select only the tools that best meet their specific solution requirements.

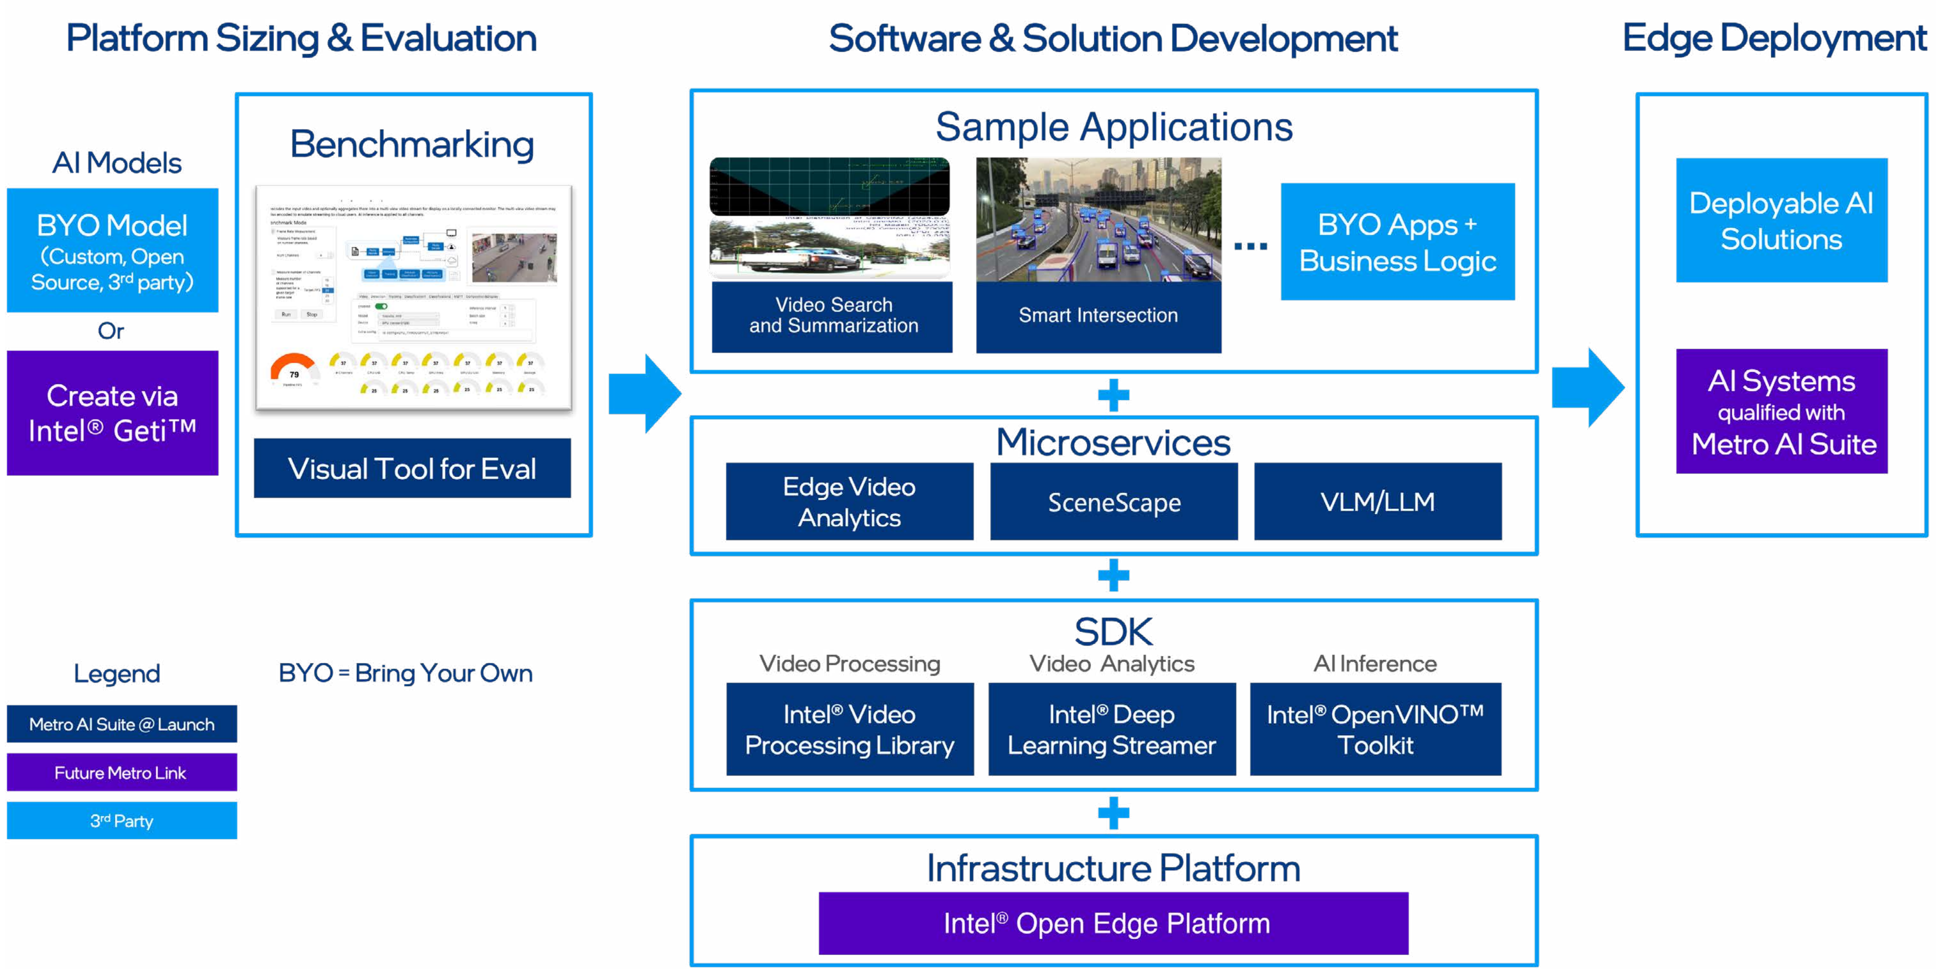

Figure 14. Metro AI Suite Software Development Kits & Visual AI Tools

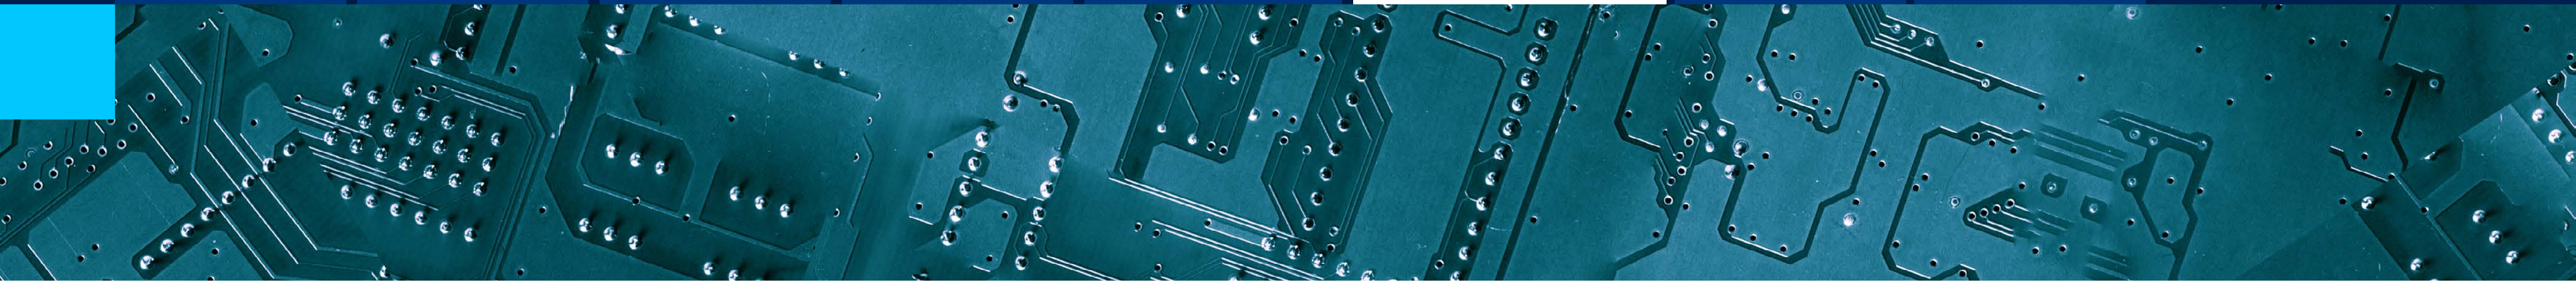

5.1 Software Development Kit (SDK)

The SDK is a comprehensive and modular toolkit for accelerated media processing and AI inference, designed to fast-track the development of visual AI solutions. It enables simple installation of Linux, kernel, graphics, Intel® Video Processing Library (Intel® VPL), Intel® OpenVINO™ toolkit, Intel® DL Streamer, etc., for quick and easy design and development of media and AI-based solutions at the edge.

Learn more about Metro AI Suite Software Development Kit and its component functionalities:

1. **Intel VPL** is a cross-platform library that offers a unified API for accessing hardware-accelerated video-decoding, -encoding, and -processing capabilities. It is designed to streamline video-processing tasks across various Intel hardware architectures, enabling developers to achieve high performance and scalability in applications such as network video recorder, multi-view display, video capture & streaming, media codec, and more.
2. **OpenVINO** is a toolkit that accelerates the development and deployment of computer vision applications by optimizing deep learning models for various Intel hardware, ensuring high performance and efficiency. It is particularly beneficial for visual AI-based Smart Roadside Unit (RSU) applications, enabling real-time processing of visual data with low latency, ideal for enhancing safety and traffic management in smart transportation systems.

   Learn more about Intel OpenVINO.
3. **Intel DL Streamer** is a framework that integrates with GStreamer to enable efficient deployment of deep learning inference pipelines for video analytics, leveraging Intel's hardware acceleration. It is ideal for visual AI-based Smart Roadside Unit (RSU) applications, as it facilitates real-time video processing for tasks like traffic monitoring and object detection, enhancing safety and traffic management with low latency.

   Learn more about Intel® DL Streamer.

5.2 Benchmarking Tool

Visual Pipeline and Platform Evaluation Tool (ViPPET) is an AI benchmarking tool that enhances both developer and user experiences through its web-based interface. It offers two main functions: visual pipeline prototyping, which provides customizable end-to-end AI and video processing proxy pipelines, and platform evaluation, which benchmarks visual AI and media performance on Edge AI-qualified hardware. This tool helps in sizing AI platforms for specific pipelines and understanding CPU and GPU utilization, assisting decision makers and developers in determining if an Intel platform suits their visual analytics needs. Ultimately, it enables TCO optimization by guiding the selection of the recommended Intel-powered hardware that leverages heterogeneous computing across CPU, iGPU, dGPU, and NPU.

Learn more about ViPPET.

5.3 Scene Analytics Tool

Intel® SceneScape is a microservice framework that streamlines advanced visual data processing and scene analysis. It offers scalable tools for interpreting complex visual environments, ideal for smart intersections and AI-driven solutions. Utilizing Intel's hardware and software, Intel SceneScape delivers real-time insights and analytics, boosting the performance of visual data applications.

Learn more about Intel SceneScape.

5.4 AI Model Training Tool

Intel® Geti™ is a versatile AI platform designed to simplify the development and deployment of computer vision models for technical system integrators and end customer ecosystems. It offers intuitive tools for data annotation, model training, and optimization, enabling seamless integration with existing systems. Leveraging Intel's robust hardware and software capabilities, Intel Geti ensures efficient and scalable AI solutions, empowering system integrators to deliver customized applications that enhance operational insights and decision-making for end customers.

Learn more about Intel Geti .

5.5 Sample Application

These are sample software applications that can be used as proof points or sample code for demonstrating the potential for implementing vertical-specific solutions on Intel-powered hardware platforms. The following section provides details on functions.

Sample Application	Function
Smart Intersection	Smart Intersection demonstrates how Edge AI technologies can address traffic management challenges using scene-based analytics.
Video Search and Summarization	Video Search and Summarization utilizes VLMs, computer vision, and audio analysis to generate concise and informative summaries of long-form videos as well as LangChain, multimodal embedding models, and agentic reasoning to enable efficient and intelligent search over video content directly at the edge.

Table 3: Metro AI Suite Sample Applications (modular, royalty-free)

5.6 Platform Blueprint

Platform Blueprint	Function
Sensor Fusion for Traffic Management	Sensor Fusion for Traffic Management includes two end-to-end visual AI pipelines, corresponding to two use cases of sensor fusion for Intelligent Traffic Management: one camera + one mmWave radar (1C+1R) and four cameras + four mmWave radars (4C+4R).

Table 4: Metro AI Suite Platform Blueprints (modular, royalty-free)

6. Demo

Scan the QR Code to watch a demo of an Intelligent Road Infrastructure Use Case

7. Appendix

7.1 Glossary

Term	Definition
Intel® Arc™ GPU	Intel's discrete graphics processing unit, such as the Intel® Arc™ A770, engineered for high-performance video analytics and AI workloads.
Intel® Deep Learning Streamer (Intel® DL Streamer)	Intel's open-source framework built on GStreamer, optimized for constructing real-time video analytics pipelines with hardware acceleration on Intel platforms, supporting multi-stream processing.
Edge Computing	A distributed computing architecture that processes data at or near its source such as RSUs minimizing latency and reducing reliance on centralized servers for time-sensitive tasks.
HEVC	High-Effiency Video Coding, a compression rate delivering 1080p video at half the bitrate of H.264 (e.g., 2 Mbps at 30 fps) while preserving visual quality, widely adopted in ITS applications.
ITMS	Intelligent Traffic Management System.
Inference Rate	The frequency of artificial intelligence model predictions per second.
IP66/67	Ingress Protection ratings per IEC 605529. IP66 ensures dust-tight enclosures and resistance to powerful water jets (100L/min); IP67 extends protection to 1-meter submersion for 30 minutes.
Metro AI Suite	Intel's end-to-end toolkit for developing, deploying, and managing AI applications across edge and server platforms, featuring prebuilt modules and workflows.
NEMA	A National Electrical Manufacturers Association standard for enclosures, providing weather and corrosion resistance, validated by tests such as 200 hours of exposure to salt spray.

Term	Definition
Natural Language Processing (NLP) Search	Natural Language Processing functionality enabling text-based queries of video data powered by advanced Vision Language Models in the Traffic Forensics Server.
OpenVINO™	Intel's open-source toolkit for optimizing deep learning models like Yolo-V5S and ResNet-50, reducing latency on Intel hardware platforms.
Sensor Fusion	The integration of data from multiple sensors (cameras, radar, lidar) into a cohesive environmental model, improving accuracy and reliability.
Traffic Forensic Server	A near-edge server deployed in Traffic Management Centers (TMCs), equipped with GenAI and a Vector Database to analyze and index up to 10TB of historical video and sensor data.
Traffic Management Center (TMC)	A centralized hub aggregating RSU data, analyzing traffic patterns, and coordinating responses across a region, typically equipped with high-performance servers and visualization tools.
Vector Database	A specialized database (e.g., Pinecone) that stores processed video lips as vectors, enabling efficient indexing and retrieval for AI-driven searches within large datasets.
VLM	Vision Language Model
Vehicle-to-Everything (V2X) Communication	Vehicle-to-Everything protocols (e.g., C-V2X, DSRC) facilitating data exchange between vehicles, RSUs, and infrastructure, enhancing coordination in real time.
YOLOv5s	A lightweight version of the You Only Look Once (v5) object detect model, optimized for edge devices, processing images at 640x640 resolution with INT8 precision.

7.2 List of Figures

7.3 List of Tables

7.4 References

Citation	Description
1	Harbor Research Roadside Solution Market Modeling (2024)
2	Beacon Edge Market Model for Edge Servers (2025)
3	Cybersecurity and Artificial Intelligence: Threats and Opportunities (2023)
4	Flashing Red Lights: Cybersecurity for Intelligent Transportation Systems (2024, United States Cybersecurity Magazine, Winter 2024 Issue)
5	Play Ransomware Group Claims Responsibility for Disrupting Kansas City Scout System (2024, The Cyber Express Magazine)

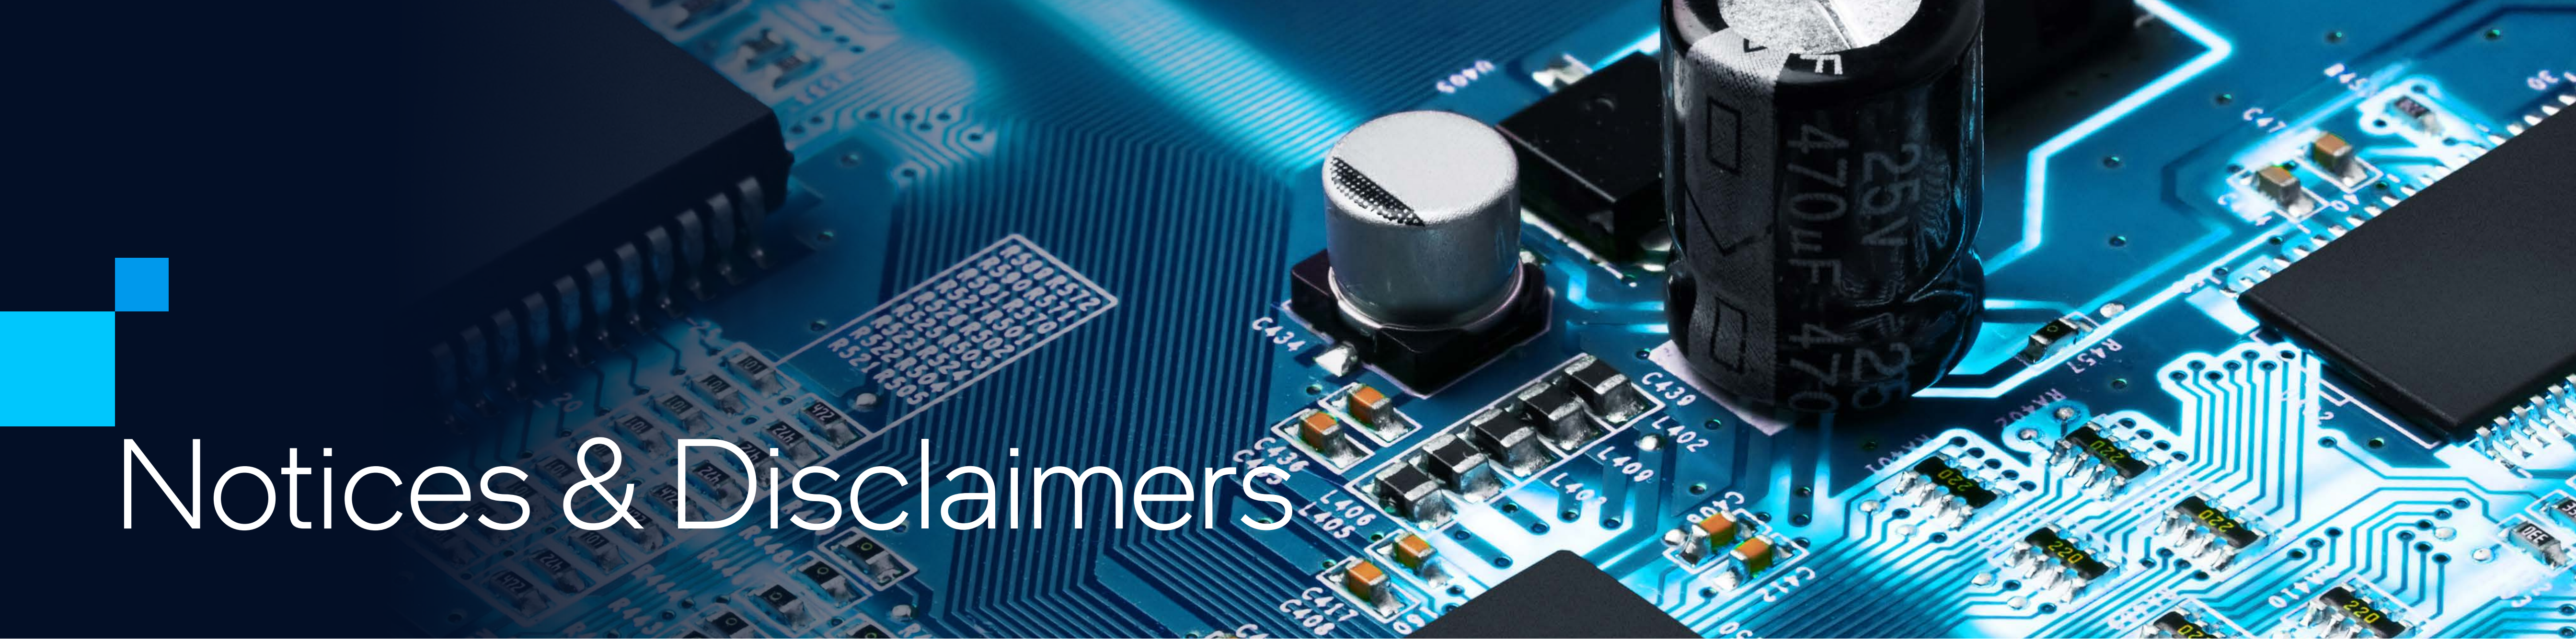

Notices & Disclaimers

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visiting the Intel Resource and Documentation Center.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex. Differences in hardware, software, or configuration will affect actual performance. Your results may vary.

Intel technologies may require enabled hardware, software, or service activation. Learn more at intel.com, or from the OEM or retailer. Performance results are based on testing as of dates reflected in the configurations and may not reflect all publicly available updates. See configuration disclosure for details.

Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See Intel's Global Human Rights Principles. Intel's 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 technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration.

No product or component can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

© 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.