Metro AI Suite Guide for Hardware Partners

Version 3.0 - June 2026 Update

Metro AI Suite for Hardware Partners

This guide is intended to help Intel Hardware partners leverage Metro AI Suite to accelerate their business

Contents

Metro AI Suite: An Ecosystem Framework to Enable Hardware Partners to Qualify and Maximize AI Performance on Intel Platforms

Hardware Catalog: AI-ready hardware systems recommendations

Metro AI Suite

Benchmarked performance across Intel® platforms

AI System qualification

- Ready to use Demos of Smart City & Transportation Sample Apps
- Benchmarking tools for evaluating Vision AI workloads
- Device qualification to ensure systems are “AI ready”

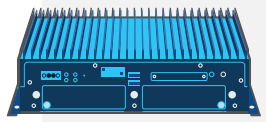

Realize AI potential for your systems

New

Use Metro AI Suite to qualify, understand, and promote system AI capabilities

Bring Your AI-Ready Intel®-Based Systems

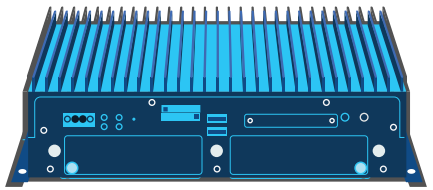

Qualify and Promote

Run Vision + AI Qualification Package

ESQ Qualification Tool

List qualified systems in Intel AI-Ready Hardware Catalog

Measure AI performance

Use VIPPET tool to easily measure AI performance (GUI tool)

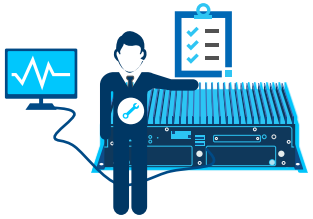

Demo AI capability

Install Visual & Gen AI Sample Applications

Demo (no-code) AI demos to customers

Qualify and Promote AI Systems

New

Intel® Edge System Qualification (Intel® ESQ) for Metro AI Suite

ESQ is a self-service qualification framework that enables hardware partners to test and validate their edge systems for AI workloads.

Tests & Measures:

- Generative AI throughput (tokens/sec)
- Multi-stream Vision AI capacity
- Media decode performance
- GPU frequency & utilization
- System memory bandwidth & latency
- Configuration and telemetry reporting

Vertical Workload Benchmarks

- Proxy pipelines such as
  - Smart NVR,
  - VSaaS Gateway,
  - Headed Visual AI,
  - License Plate Recognition
- Light - Medium - Heavy AI Workloads

New

Qualify your Intel Edge AI Systems with ESQ

Use Metro AI Suite Device Qualification Tool

- **Intel® Edge System Qualification (Intel® ESQ) for Metro AI Suite** allows customers to run an Intel-provided test suite on the target system, enabling partners to qualify their platform as an Edge AI ready device
- See link for instructions, supported system requirements, & other prerequisites

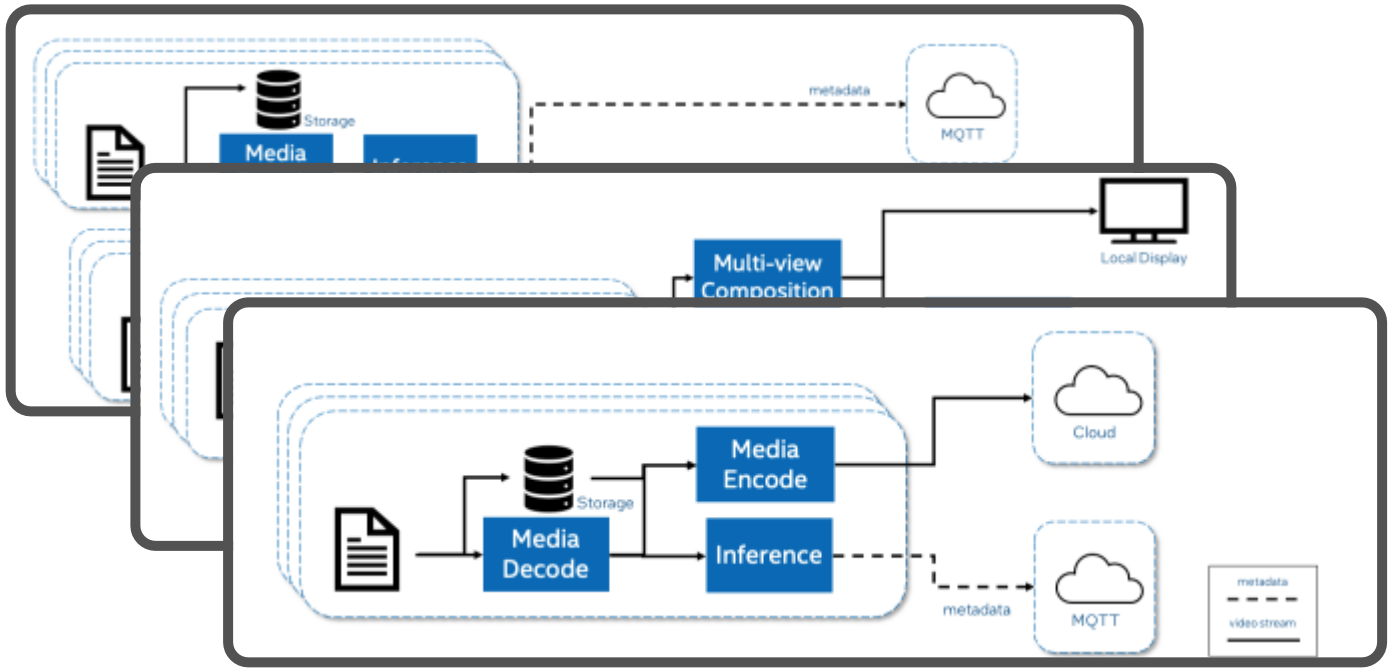

Why Qualify Systems?

- Ensure AI-ready systems for ISV and ecosystem partners
- Enable Hardware Catalog eligibility and ecosystem visibility
- Remove guesswork with standardized, data-driven system qualification
- Support confident system positioning for Vision AI and Generative AI use cases
- Access performance benchmarks across AI, media, memory, and GPU behavior

Eligible Intel platforms

- Intel® Core™ Ultra Series 3
- Intel® Core™ Ultra Series 2
- Intel® Core™ Series 2
- Intel® Processor for Desktop
- Intel® Processor X-series
- Intel® Processor N-series
- Intel® Xeon® 6 Processors
- 5th Gen Intel® Xeon® Scalable Processors
- Intel® Xeon® W Processors
- Discrete GPU Options
  - Intel® Arc™ B-Series Graphics
  - Intel® Arc™ Pro B-Series Graphics

New

Partner Self-Qualification using Intel® ESQ

Four Easy steps for Partner Self Qualification

Download & Install Intel AI Edge System ESQ framework	Run ESQ Profile (Test Modules)	Submit Results Summary for Self-Qual	Approved and Listed
1  Download ESQ from GitHub	2  Run ESQ Tests	3  Submit Application	4  Approval and Listed

OXM | Intel

intel ai Eligible for Partner Endorsement Badge & additional partner benefits.

The Recommended Hardware Catalog helps SW partners find qualified edge systems

New

Submit Qualified Systems to Hardware Catalog

Metro AI Suite – Recommended Hardware Catalog

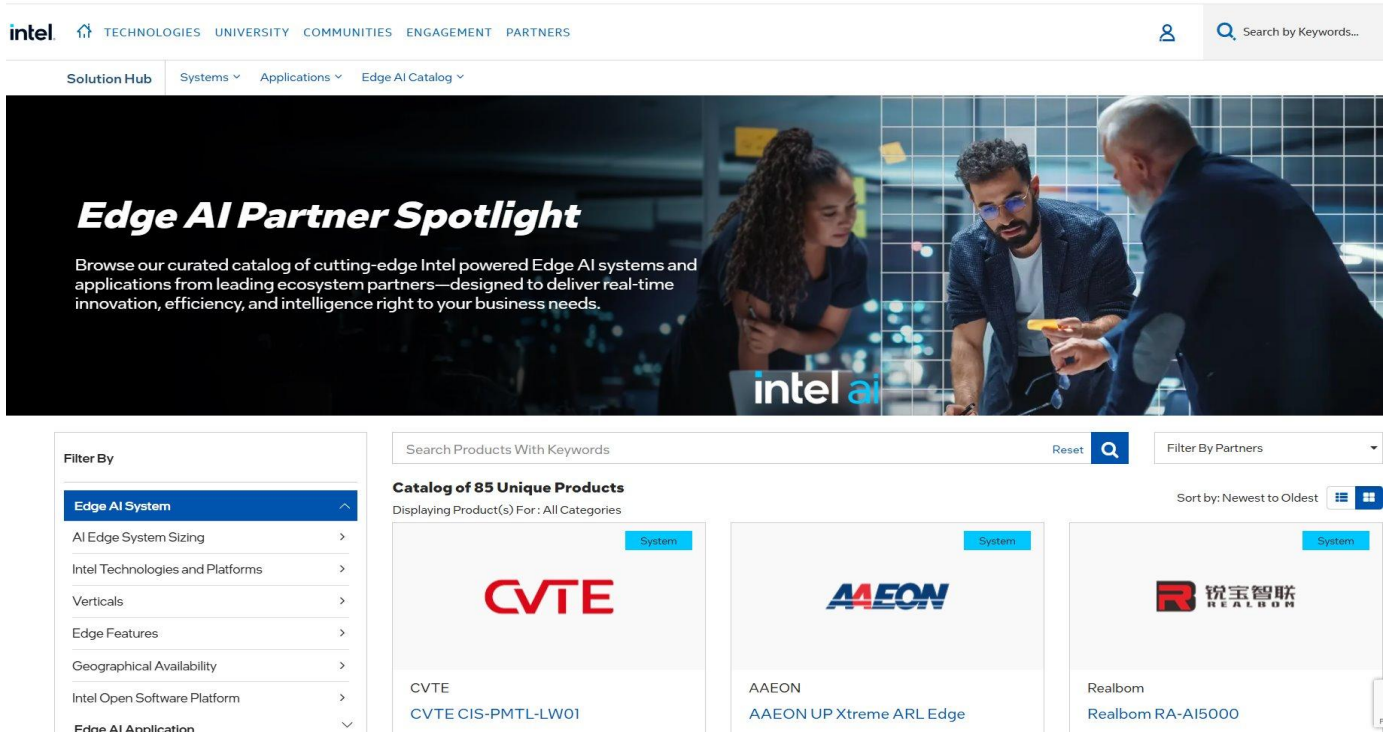

Why add your systems to the Catalog?

- Build awareness of your AI system portfolio
  - Available to ISVs, SIs, and End Users looking for AI systems
- Users can filter systems by processor type, industry, product application, geography, and more
- Only Metro AI Suite-qualified hardware listed

To upload, follow instructions after submitting ESQ results

System Benchmarking with Real-world Workloads

ViPPET (Visual Pipeline and Platform Evaluation Tool)

Choose the right platform with real pipeline evidence

1
Choose base pipeline

Select a video workflow

2
Customize pipeline

Choose models and other parameters to match your workflow

3
Optimize & Benchmark

ASSIGN WORKLOADS TO

CPU iGPU NPU

Measure throughput, latency, utilization & TCO

Flexible stage-to-hardware mapping reflects real deployment choices.

- See how many camera channels run on CPU, iGPU, and NPU.
- Reduce rollout risk with evidence-based hardware choices.
- Benchmark end-to-end media and AI workloads on real pipelines.
- Auto-optimize for the best throughput, power, and TCO.

Explore details →

Easy End-to-End Visual AI Benchmarking with ViPPET

Why ViPPET?

Create a new pipeline or select pre-defined pipelines

Decision-Ready Insights

- Evaluate real-world Visual AI performance on Intel hardware
- Measure FPS & device utilization across full video + AI pipelines

Faster System Evaluation

- Quickly see how your workload scales across CPU / GPU / NPU.
- Compare platforms to choose the right hardware with confidence.

Performance Optimization

- Experiment with models and parameters
- Map individual pipeline stages — or the whole chain — to CPU, GPU, or NPU
- Identify the highest-throughput configuration for your solution with the built-in optimizer.

Predictable Deployment

- Use proxy pipelines that mirror production workloads.
- Reduce guesswork and optimize TCO early in development.

Explore details →

Easy End-to-End Visual AI Benchmarking with ViPPET

Getting Started 🚀

Select a ready-made pipeline and run it on your device in just a few clicks to assess your Intel hardware's end-to-end Visual AI capability.

Advanced Usage

Select a ready-made pipeline or bring your own using a DLStreamer pipeline string

Configure the pipeline by choosing or swapping proxy models, videos, devices, and adjusting step-specific parameters such as batch size, parallel inference requests, inference interval for different stages.

Rerun and compare pipelines to benchmark detailed end-to-end (E2E) performance across different configurations.

Track and analyze results in the Jobs section to understand system capacity for Visual AI workloads.

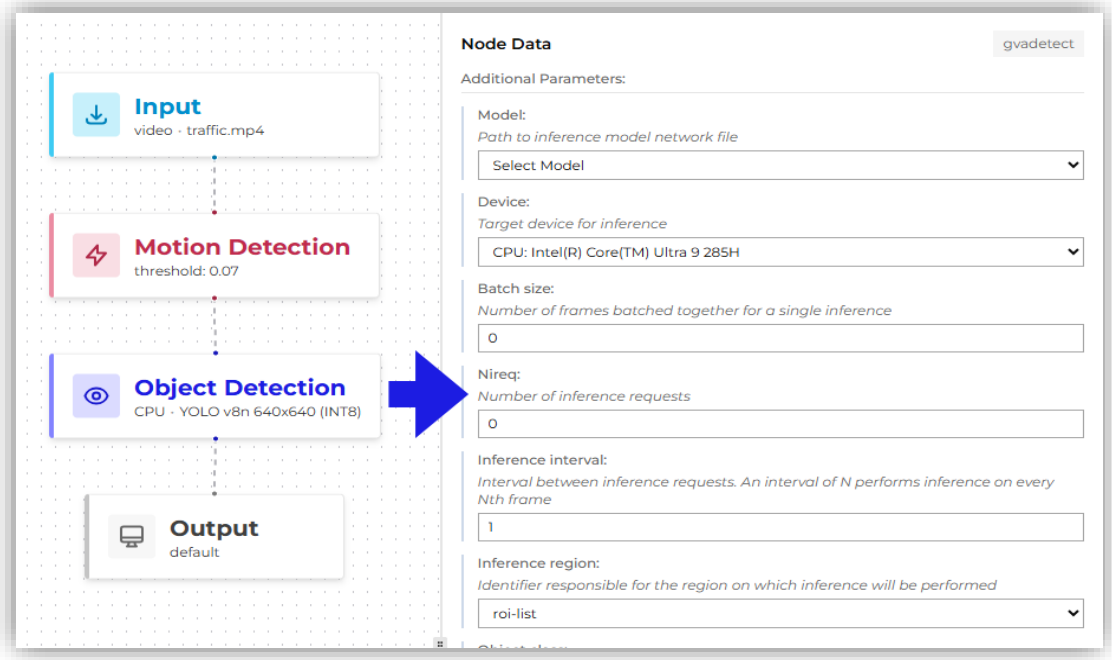

Pipeline Customization

Explore details →

Easy End-to-End Visual AI Benchmarking with ViPPET

Benchmark end-to-end media and AI workloads on real pipelines.

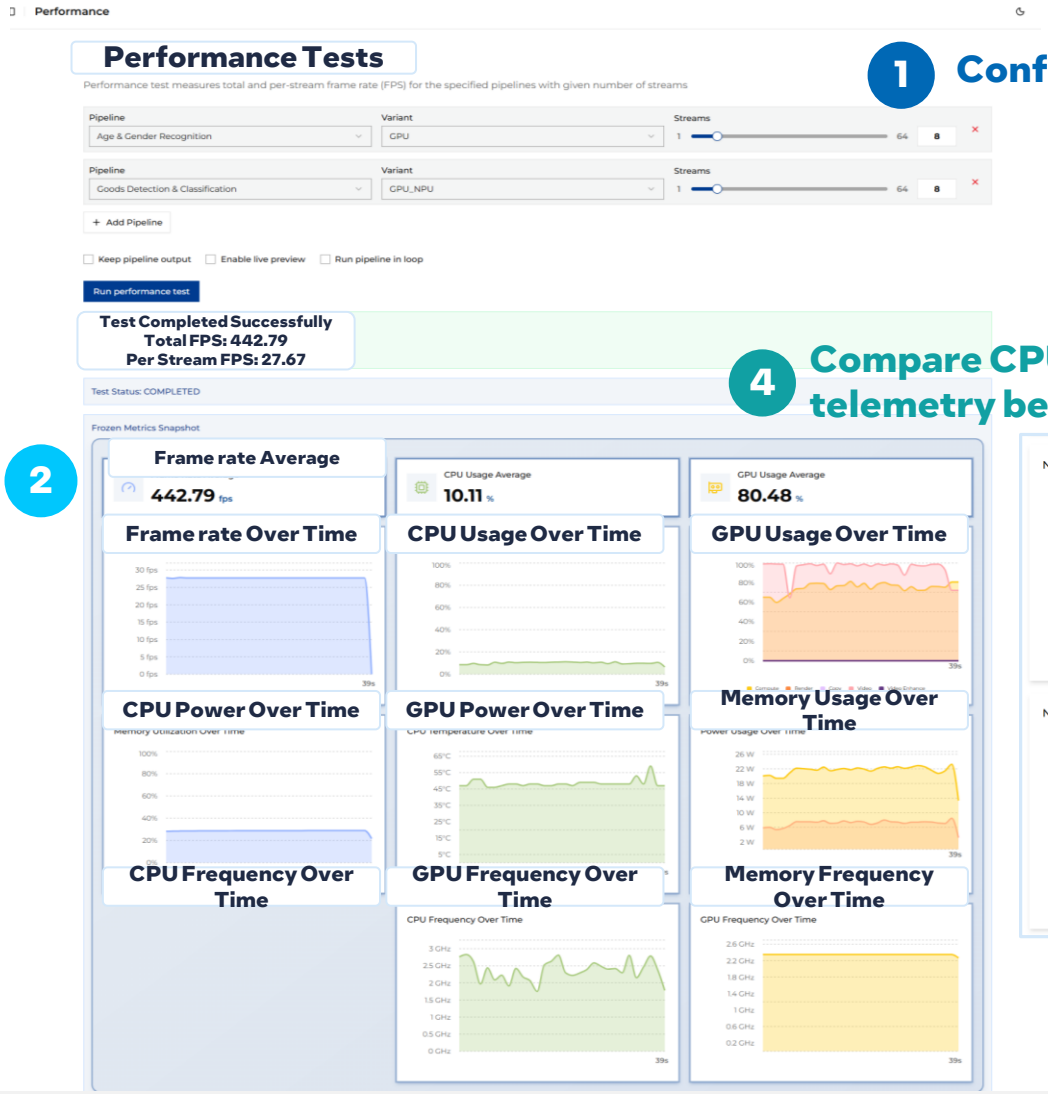

1 **Configure a real pipeline test**

Confirm the workload has target FPS. 2

Monitor live latency and spot latency spikes 3

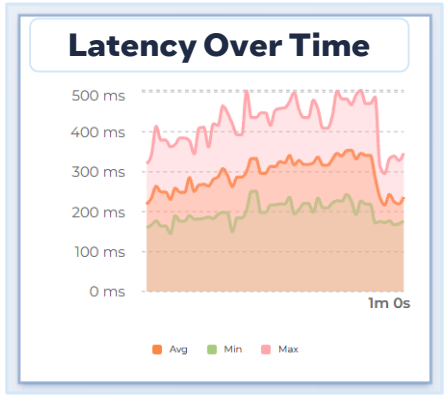

4 **Compare CPU, iGPU and now NPU telemetry before deployment.**

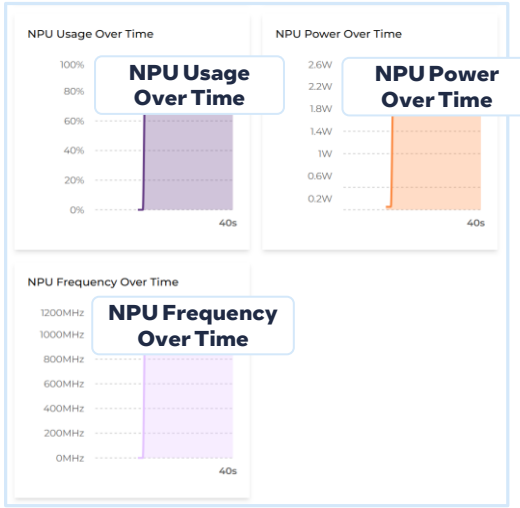

Explore details →

New

Easy End-to-End Visual AI Benchmarking with ViPPET

Flexible Pipelines, Deeper Performance Visibility

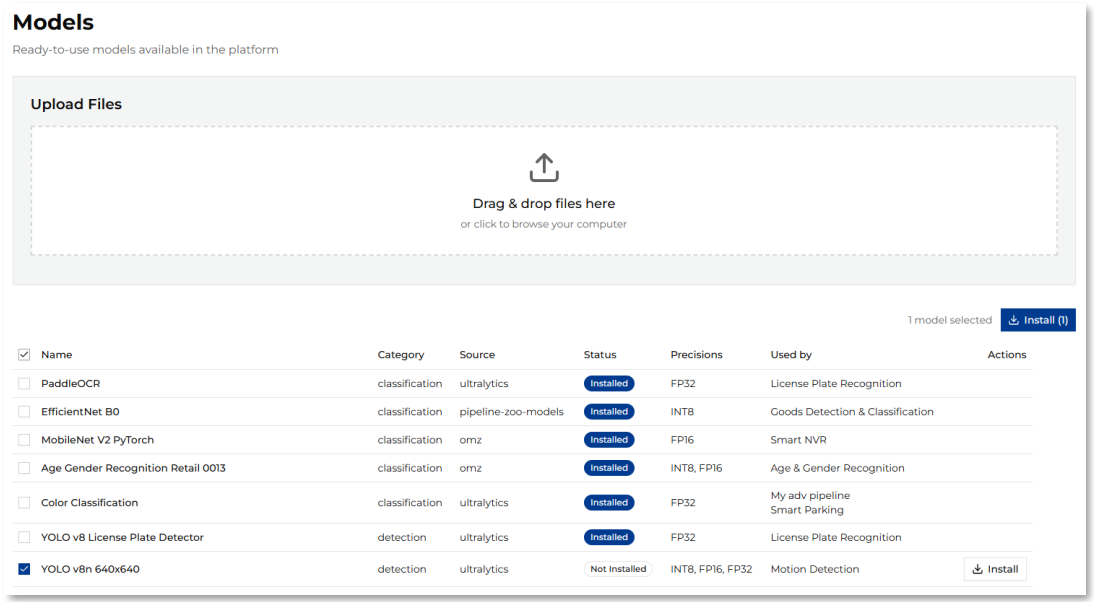

New Predefined Pipelines supporting VLMs for more vertical business use cases.

Custom OpenVINO and Intel Geti model uploads for easy bring-your-own-model testing and reuse.

Camera input, video and image-set uploads for flexible, real-world pipeline inputs.

Real-time pipeline output for improving experience.

NPU utilization reporting alongside CPU and GPU for full visibility into AI pipeline performance on modern Intel hardware.

End-to-end latency metrics alongside throughput for spotting bottlenecks and verifying real-time pipeline performance.

Explore details →

Showcase AI Capabilities

Run Visual AI and Gen AI sample applications for demo purposes

Reference Visual and Gen AI Sample Apps, Tutorials and Blueprints

Why Use AI Sample Apps?

- Understand & evaluate Intel platforms for Computer Vision & Gen AI use cases
- Helps developers streamline code or jumpstart development

Sample Applications, Tutorials, Blueprints

- ✓ Smart Parking
- ✓ Smart Intersection
- ✓ Loitering Detection
- ✓ Smart Tolling
- ✓ Reidentification (Image Based Video Search)
- ✓ Interactive Digital Avatar
- ✓ Gen AI Based Video Q&A
- ✓ Sensor Fusion for Traffic Management
- ✓ Video Processing for NVR
- ✓ Video Search and Summarization
- ✓ Live Video Applications
- ✓ NVIDIA → Intel Migration Guide
- ✓ Live Video Alert Agent
- ✓ Smart Traffic Intersection Agent
- ✓ Smart Route Planning Agent
- ✓ Smart NVR
- ✓ Agentic AI Tutorials
- ✓ Blueprints
  - Vision-AI Based Agentic Predictive Maintenance
  - Multimodal Predictive Maintenance
  - Smart Buildings Digital Twin Blueprint

Smart Parking

AI-driven parking management application that analyzes live video with pre-trained deep learning models to detect, count, and track open spaces and report parking availability in real time.

Real-time parking space detection and occupancy monitoring

Key Features:

- **Modular, microservice-based architecture** (including Intel® DL Streamer)
- **Vision Analytics Pipeline:** Detect and classify objects using pre-configured AI models. Customize parameters (thresholds and object types) without requiring additional coding.
- **Integration with MQTT, Node-RED, and Grafana:** Facilitates efficient message handling, real-time monitoring, and insightful data visualization.
- **User-Friendly:** Simplifies configuration and operation through prebuilt scripts and configuration files.
- **Made with Visual AI Demo Kit** low-code app framework

Supported Intel Platforms:

- Intel® Core™, Intel® Core™ Ultra, Intel® Xeon® platforms

Explore details →

Smart Intersection

Edge-AI traffic management application that fuses multi-camera, scene-based analytics — including lidar and radar — to track vehicles and pedestrians across an intersection and deliver real-time, coordinated traffic insights.

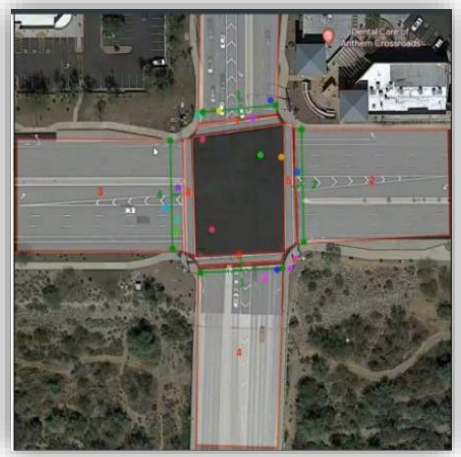

Real-time, multi-sensor traffic insights across a complex intersection

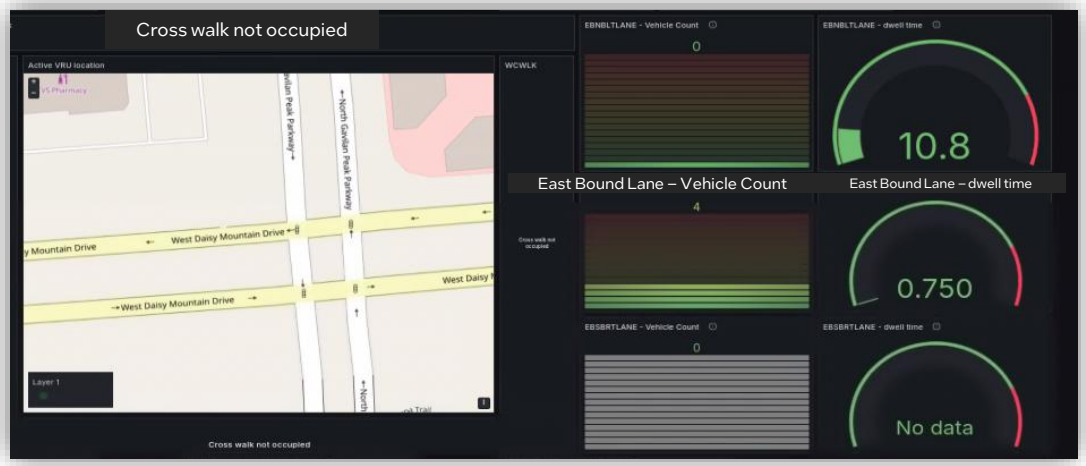

Key Features:

- **Multi-sensor integration, including cameras, lidar, and radar** help serve use cases like pedestrian safety and traffic analytics
- **Scene-based / Unified analytics:** Define regions of interest via an independent map view, simplifying multi-object tracking, motion vector analysis, and business logic across sensors
- **Integration with MQTT, InfluxDB, Node-Red, and Grafana:** Facilitates efficient message handling, near real-time monitoring, and insightful data visualization.
- **Modular, microservice-based architecture** (including DLStreamer) enables composability and reconfiguration

Supported Intel Platforms:

- Intel® Core™, Intel® Core™ Ultra, Intel® Xeon® platforms

Explore details →

Loitering Detection

AI vision application that tracks people in live video and applies dwell-time thresholds to flag lingering individuals and raise real-time alerts.

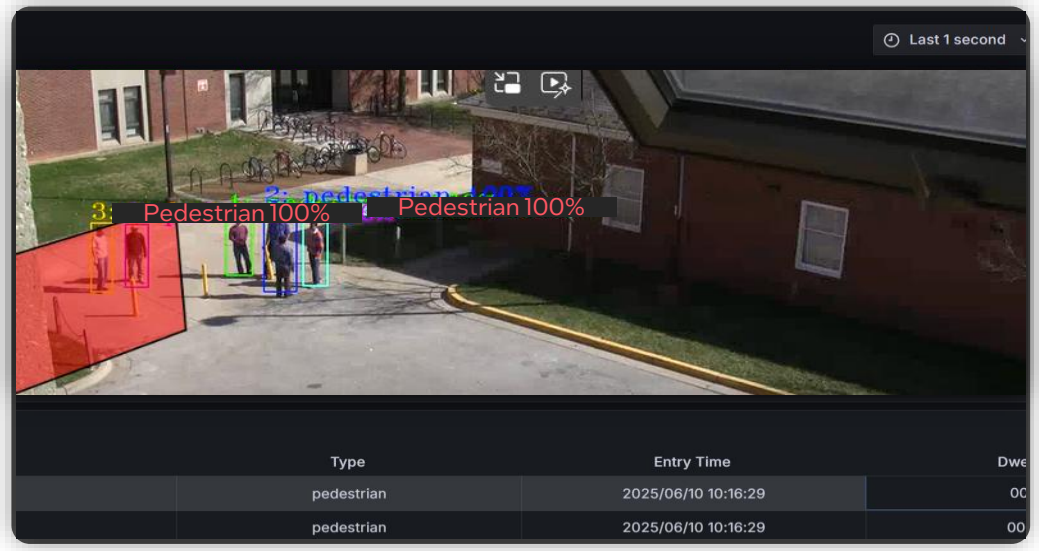

Loitering detection in action

Key Features:

- **Vision Analytics Pipeline: Detect and classify objects using pre-configured AI models.** Customize parameters such as thresholds and object types without requiring additional coding.
- **Integration with WebRTC Server, MQTT, Node-RED, and Grafana**: Facilitates efficient message handling, real-time monitoring, and insightful data visualization.
- **User-Friendly**: Simplifies configuration and operation through prebuilt scripts and configuration files.
- **Made with Visual AI Demo Kit** low-code app framework

Supported Intel Platforms:

- Intel® Core™, Intel® Core™ Ultra, Intel® Xeon® platforms

Explore details →

Smart Tolling

Edge AI application that automates vehicle identification, classification, and accurate toll calculation across multi-lane toll points.

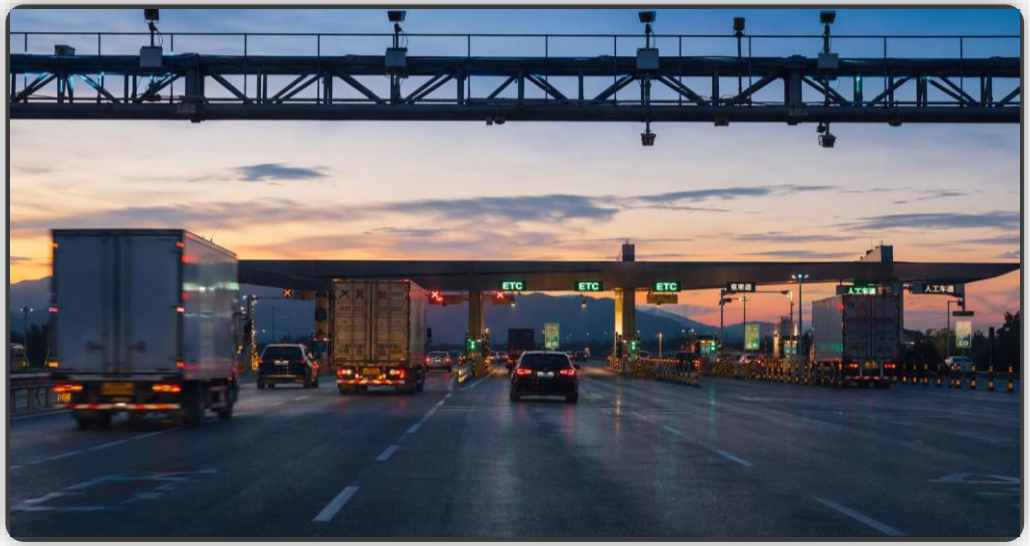

Key Features:

- **Multi-Camera Vehicle Recognition:** Reads plates and identifies vehicle type and color in real time.
- **Sensor Fusion with Intel® SceneScape**: Fuses camera views to track each vehicle through the toll zone.
- **Accurate Axle Counting**: Detects lifted axles for fair taxation
- **Operator Dashboard:** Grafana dashboard for traffic, flow, and audits on InfluxDB time-series data.

Supported Intel Platforms:

- Intel® Core™, Intel® Core™ Ultra, Intel® Xeon® platforms

Explore details →

Reidentification (Image based Video Search)

Image-based video search application that combines object detection, feature extraction, and vector search to locate a query object across live or recorded camera feeds in near real time.

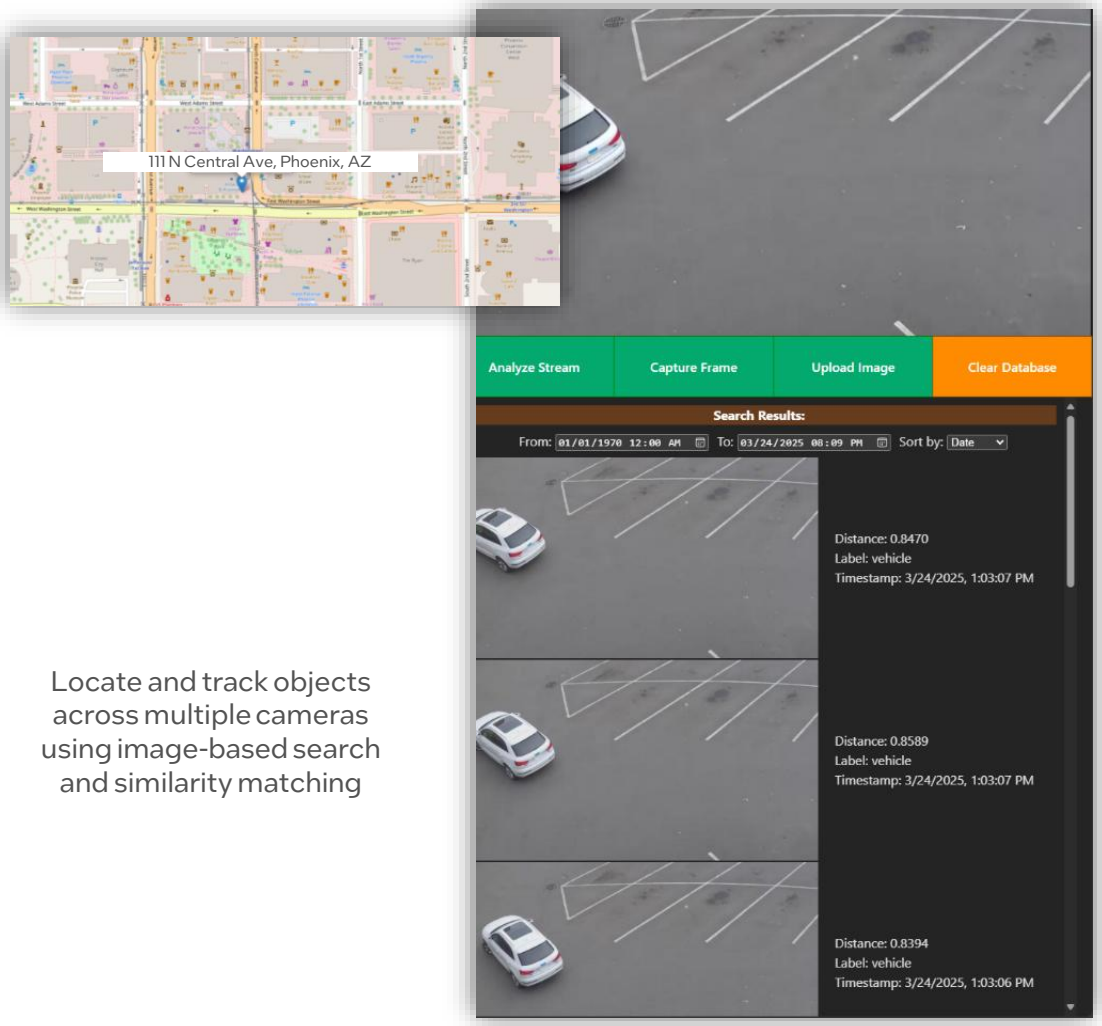

Locate and track objects across multiple cameras using image-based search and similarity matching

Key Features:

- **Enables cross-camera tracking / re-identification** – useful in both real-time and forensic investigations
- **Shows how to combine edge AI microservices** for video ingestion, object detection, feature extraction, and vector-based search.
- **Integration with DLStreamer Pipeline Server, MediaMTX, MQTT, MilvusDB, ImageIngestor**

Supported Intel Platforms:

- Intel® Core™ and Intel® Core™ Ultra platforms

Explore details →

Interactive Digital Avatar

An AI-powered avatar that understands speech, analyzes video context, and responds in real time—enabling natural, interactive experiences for users.

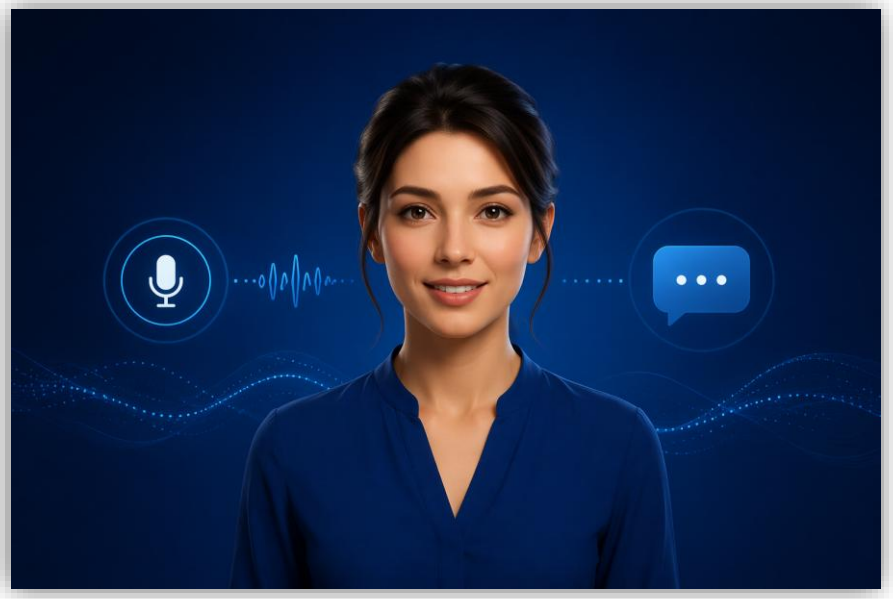

Key Features:

- **Multimodal Interaction** Combines voice, vision, and language models to understand user input and respond naturally.
- **Real-Time Conversational AI** Delivers low-latency, context-aware responses using live video, captions, and RAG pipelines.

Supported Intel Platforms:

- Intel® Arc™ A770 graphics

Explore details →

Gen AI Based Video Q&A (VQ&A)

Generative-AI assistant that uses a large vision language model to answer open-ended, natural-language questions about image and video content in real time.

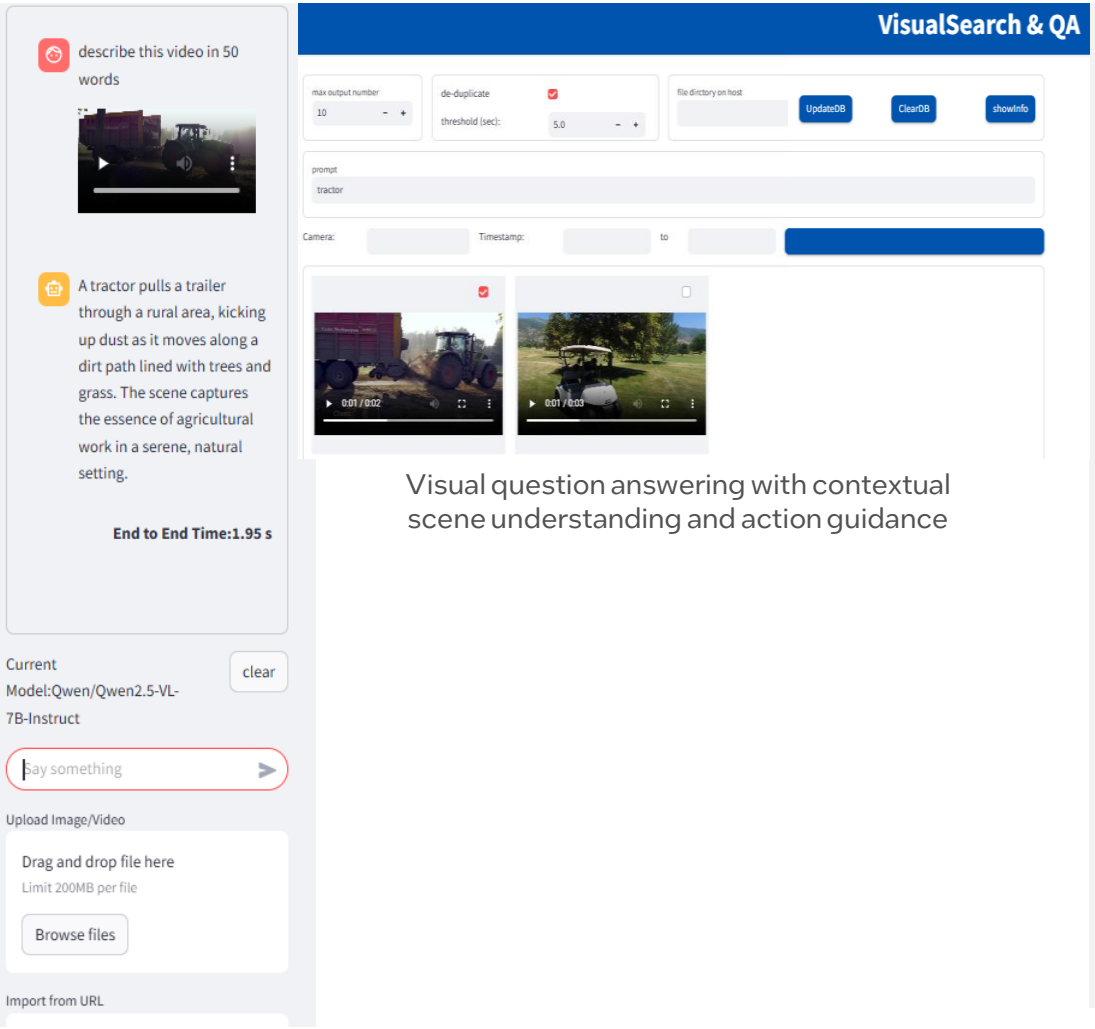

Visual question answering with contextual scene understanding and action guidance

Key Features:

- **VLM and LLM GenAI**, including video summarization
- **Integrates LVM Server, IPEX-LLM Server, & Web UI with chat**
- Supported Models: **Llava-1.5-7b, Qwen2-VL-7B-Instruct, InternVL2-4B**

Supported Intel Platforms:

- Intel® Core™ and Intel® Arc™ A-series Graphics, such as A770 GPU

Explore details →

Sensor Fusion for Traffic Management

A multi-modal sample application to accurately monitor traffic conditions by fusing camera and sensor inputs

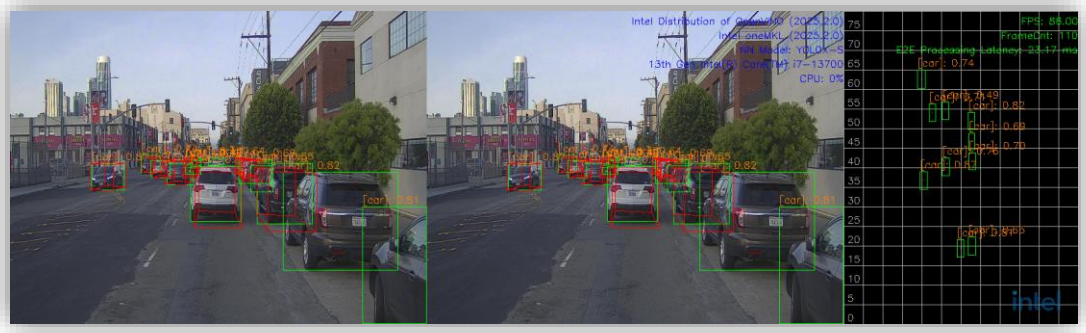

2 camera streams and 1 LiDAR sensor configuration

The sample application combines camera + LiDAR outputs to improve accuracy

Key Features:

- **Multi-modal sensor fusion:** Combine camera, radar, and lidar inputs for robust detection in fog, rain, or darkness
- **Accurate measurement:** Deliver reliable speed and distance estimation for traffic monitoring
- **End-to-end pipeline:** Lidar signal processing, video analytics, camera-lidar fusion, and visualization on a single Intel SoC
- **Scalable configurations:** Support multiple camera + lidar setups, from 1 camera +1 lidar sensor up to 12 cameras+4 Lidar censors

Supported Intel Platforms:

- Intel® Core™ Ultra processors with integrated GPU or Intel® Core™ processors with Intel® Arc™ B-Series Graphics

Explore details →

Video Processing for NVR

A sample application to evaluate and optimize NVR video processing workflows.

Key Features:

- **Concurrent multi-channel processing:** Decode, process, and display many H.264/H.265 streams at once on Intel integrated GPU
- **Video analytics & transcoding:** Run analytics and transcoding with ready-made example apps
- **Configuration-driven workloads:** Set channels, scaling, and layout via text config—no code changes
- **Multiview monitoring:** Check performance and debug pipelines across multiple 4K displays
- **Built for NVR evaluation:** Benchmark and debug core workloads in a real NVR setup

Supported Intel Platforms:

- Intel® Core™ processors with integrated GPU or Intel® Arc™ Graphics

Explore details →

Video Search and Summarization

Generative-AI video search and summarization application that uses multimodal embeddings, vector search, vision language models, and audio analysis to retrieve relevant clips from natural-language queries and condense long video streams into concise summaries.

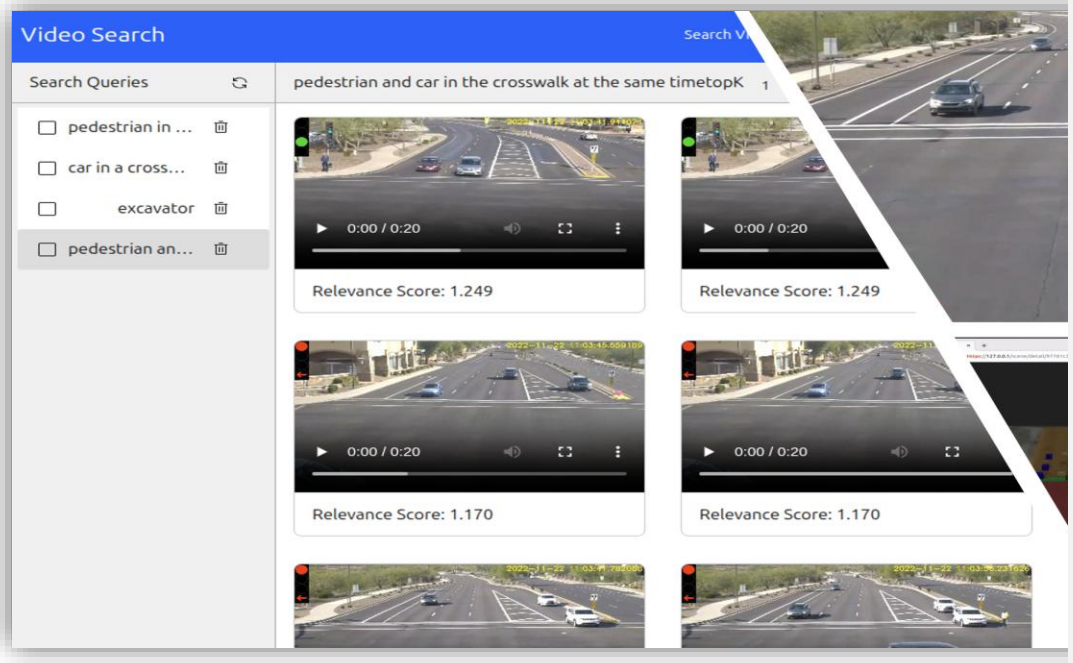

Key Features:

- **Video Search**: This functionality leverages LangChain, multimodal embedding models, and agentic reasoning to enable efficient and intelligent search over video content directly at the edge.
- **Video Summarization**: Using Vision Language Models (VLMs), Computer Vision, and Audio Analysis, the application distills key information into brief synopses from large volumes of data within long-form videos.

Supported Intel Platforms:

- Intel® Core™, Intel® Core™ Ultra, Intel® Xeon® platforms, Intel® Arc™ Graphics

Explore details →

Live Video Applications

Real-time, AI-powered analytics for live video streams on Intel platforms.

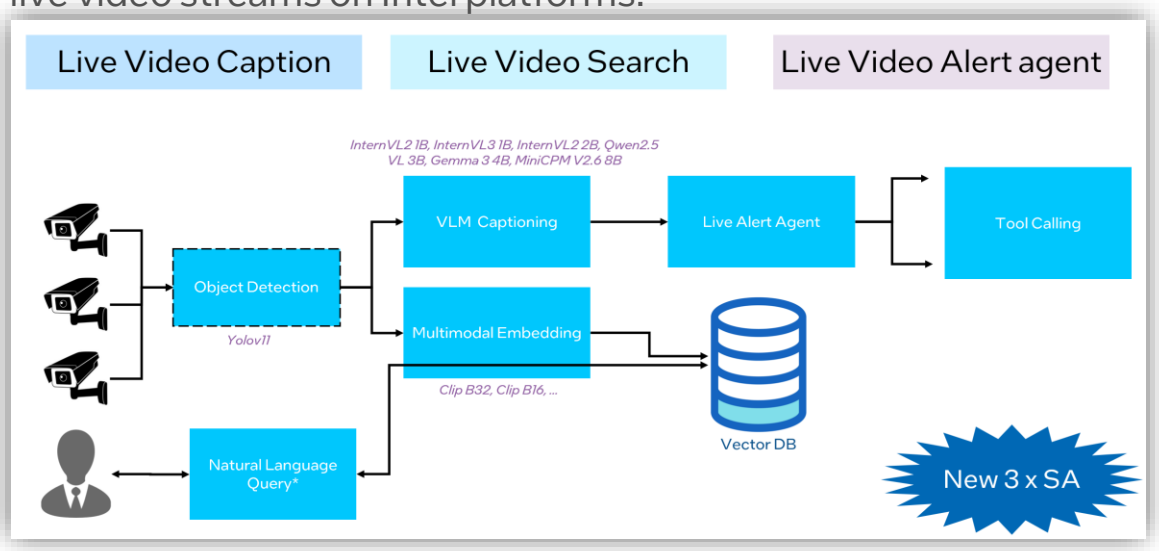

Live Video Captioning

Generates real-time, AI-powered captions from live video streams to improve accessibility and situational awareness.

Supported Intel Platforms:

- Intel® CoreTM, Intel® CoreTM Ultra processors

Explore details →

Live Video Search

Enables natural-language search across live and recorded video streams using AI-generated metadata and embeddings to quickly find relevant moments.

Supported Intel Platforms:

- Intel® Xeon® processors, Intel® Arc™ Graphics

Explore details →

Live Video Alert Agent

New

Monitors live video streams in real time to detect events and trigger configurable AI-driven alerts for rapid response.

Supported Intel Platforms:

- Intel® Core™, Intel® Xeon® processors

Explore details →

Live Video Captioning RAG

Converts live video captions into a RAG-powered knowledge base for natural-language search and contextual insights.

Supported Intel Platforms:

- Intel® CoreTM, Intel® CoreTM Ultra processors

Explore details →

NVIDIA → Intel Migration Guide

Structured migration of AI pipelines — across models, pipelines, and serving

	NVIDIA STACK		INTEL STACK
Convert models *Translate to OpenVINO-supported formats*	Native framework formats	→	OpenVINO IR (.xml / .bin)
Replace pipelines *Migrate media and AI graphs to DL Streamer*	DeepStream	→	DL Streamer
Swap frameworks *Move from the CUDA stack to OpenVINO*	CUDA-based stack	→	OpenVINO stack
Update serving *Drop-in, API-compatible inference server*	Triton Inference Server	→	OpenVINO Model Server
Deploy on Intel HW *Unified runtime across CPU, iGPU, and NPU*	GPU only	→	CPU + iGPU + NPU (unified)

Same workloads · broader hardware reach · lower TCO

Explore details →

Live Video Alert Agent

AI-powered video alerting for RTSP streams using OpenVINO™ Vision Language Models and natural-language prompts

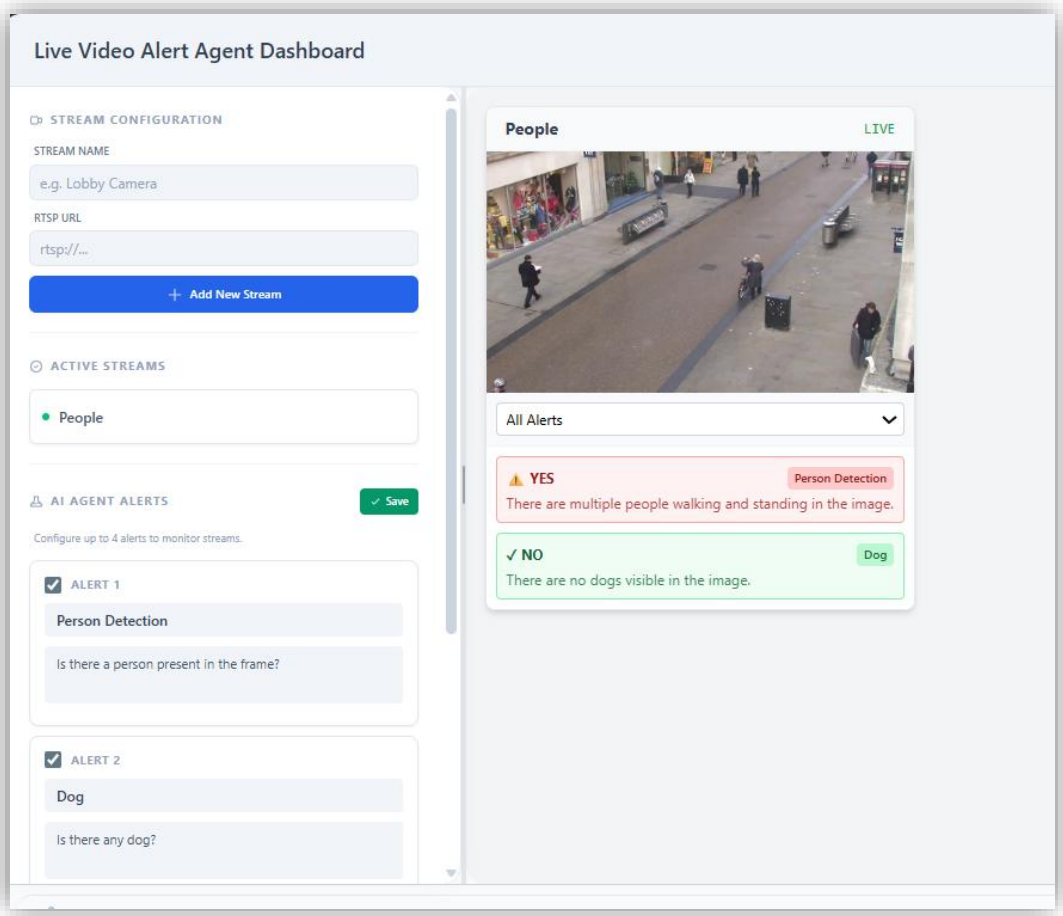

RTSP camera streams are analyzed in real time with user-defined alert prompts and low-latency dashboard events.

Key Features:

- **Dynamic alert prompts:** Define and modify alerts in real time through the UI without redeploying.
- **Real-time event broadcasting:** Alerts are displayed to a unified dashboard with low latency.
- **Modular architecture:** Decoupled ingestion, VLM analysis, and event broadcasting enable scalable deployment.
- **Telemetry:** Dashboard tracks CPU, GPU, and memory utilization, with health/readiness/metrics endpoints for operational visibility.

Supported Intel Platforms:

- Intel® Core™ Ultra 2 and 3 processors with integrated GPU, Intel® Xeon® platforms with/without Intel® Arc™ GPU

Explore details →

Smart Traffic Intersection Agent

Edge AI service for real-time intersection monitoring, directional traffic density, and VLM-powered traffic insights

Intersection cameras and traffic data streams are analyzed at the edge to estimate density, surface alerts, and support driving recommendations.

Key Features:

- **Real-time intersection analytics:** Monitor traffic conditions and directional vehicle density at the edge.
- **VLM-powered insights:** Use vision language models to interpret scene context and traffic scenarios.
- **Traffic data streaming:** MQTT broker support streams traffic data; camera images are managed for downstream services.
- **Open interfaces:** REST APIs and a UI help configure workflows, review outputs, and integrate other agents.

Supported Intel Platforms:

- Intel® Core™ Ultra 2 and 3 processors with integrated GPU or Intel® Xeon® processor

Explore details →

Smart Route Planning Agent

AI-powered route optimization agent that uses multi-agent communication to analyze traffic intersections and find incident-free paths between source and destination in real-time.

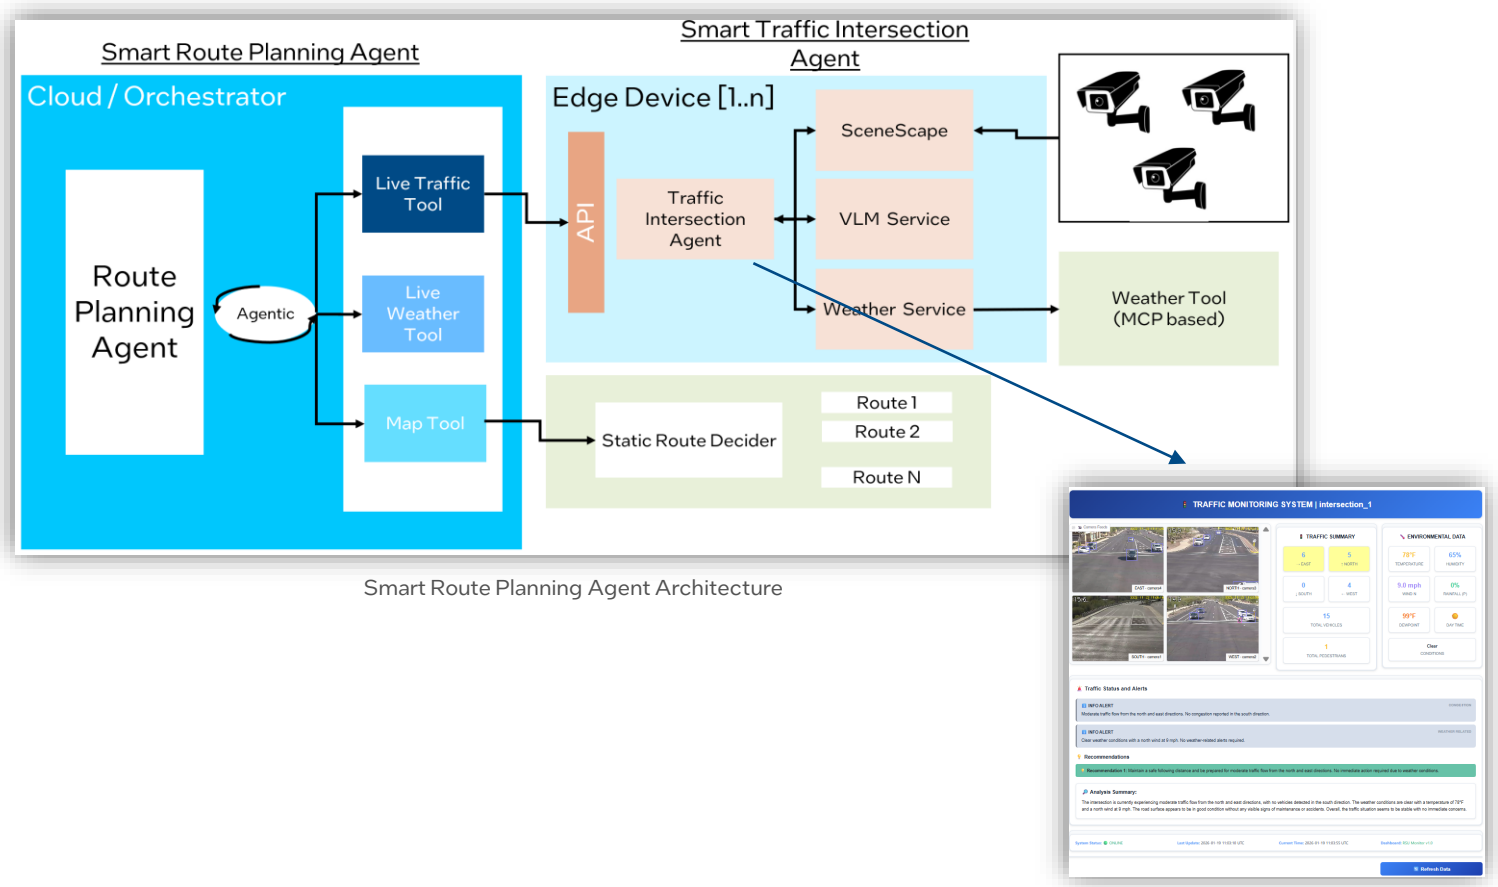

Smart Route Planning Agent Architecture

Key Features:

- **Smart Route Planning Agent (Cloud)** orchestrates edge agents and recommends optimal routes using live traffic data
- **Smart Traffic Intersection Agent(Edge)** analyzes traffic patterns, counts vehicles, and identifies congestion causes using SceneScape + VLMs
- **Integrated Tools:** Weather services, live traffic data, and mapping for comprehensive decisions

Supported Intel Platforms:

- Intel® Xeon® 6 processors

Explore details →

Smart NVR

Transforms traditional NVRs into intelligent, context-aware systems using GenAI-powered vision analytics to process video streams at the edge for real-time insights and automation.

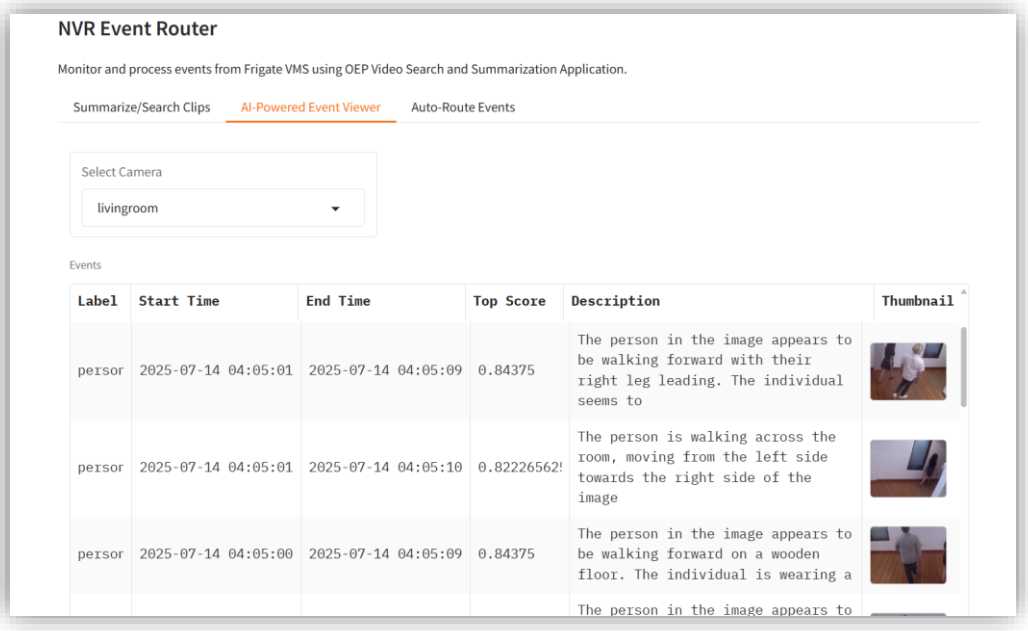

Quickly search, summarize, and understand video streams using AI-powered video analytics.

Key Features:

- **Edge-optimized analytics,** Process video locally, reducing bandwidth and enabling faster response
- **AI-powered summaries,** Generate context-aware event descriptions using vision-language models
- **Advanced video search,** Find and access key video moments with smart embedding-based search.
- **Real-time automated routing,** Automatically process and route clips using custom rules
- Supported Models: **Llava-1.5-7b, Qwen2-VL-7B-Instruct, InternVL2-4B**

Supported Intel Platforms:

- Intel® CoreTM, Intel® Xeon® and Intel® ArcTM A-series Graphics, such as A770 GPU

Explore details →

New

Agentic AI Tutorials

Learn how to design and deploy production-grade Agentic Visual AI on Intel edge hardware with step-by-step guidance and ready-to-run app recipes.

Application Security Enablement in Smart Intersection

Ensure data privacy

- Enable built-in security features with strong encryption protocols and blockchain.
- Safeguard infrastructure with advanced threat detection
- Mitigate vulnerabilities across distributed edge devices.

Explore details →

AI Crowd Analytics

Customize for metro use cases

- Use vibe coding with Metro Vision AI App Recipe
- Use pre-configured AI models to detect and classify objects
- Develop customized crowd analytics with customized parameters

Explore details →

New

Solution Blueprints

Intel-validated software blueprints enabling Agentic Predictive Maintenance for critical infrastructure and Digital Twins for smart buildings

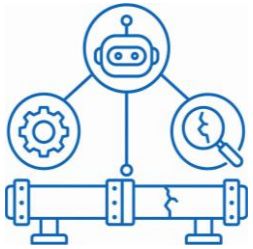

Vision-AI Based Agentic Predictive Maintenance

Multi-agent reasoning for failure detection, prediction, policy enforcement, and inspection audits.

- Vision inference, multi-agents that can be deployed in one Edge Platform or across distributed edge deployment
- Detects six pipeline defect classes in real time
- Interactive chat for report generation and audit compliance check
- Uses Public Open Data Set

Explore details →

Multimodal Predictive Maintenance

Fuses edge vision with chemical-sensor data at inference time for more reliable detection in safety-critical environments.

- Real-time multi-modal image and sensors data fusion
- Cuts false positives under ambiguous conditions
- Degrades gracefully when a camera or sensor fails
- Uses Public Open Data Set

Explore details →

Smart Building Digital Twin

Intel SceneScape unifies multi-camera vision into a live 4D digital twin where agents reason on real-world scene data, not pixels.

- Tracks every person, object, and door with metric position and history
- Powers access control, occupancy, and safety beyond video analytics

Explore details →

Vision-AI Based Agentic Predictive Maintenance

An end-to-end pipeline powered by distributed AI agents for infrastructure monitoring and inspection.

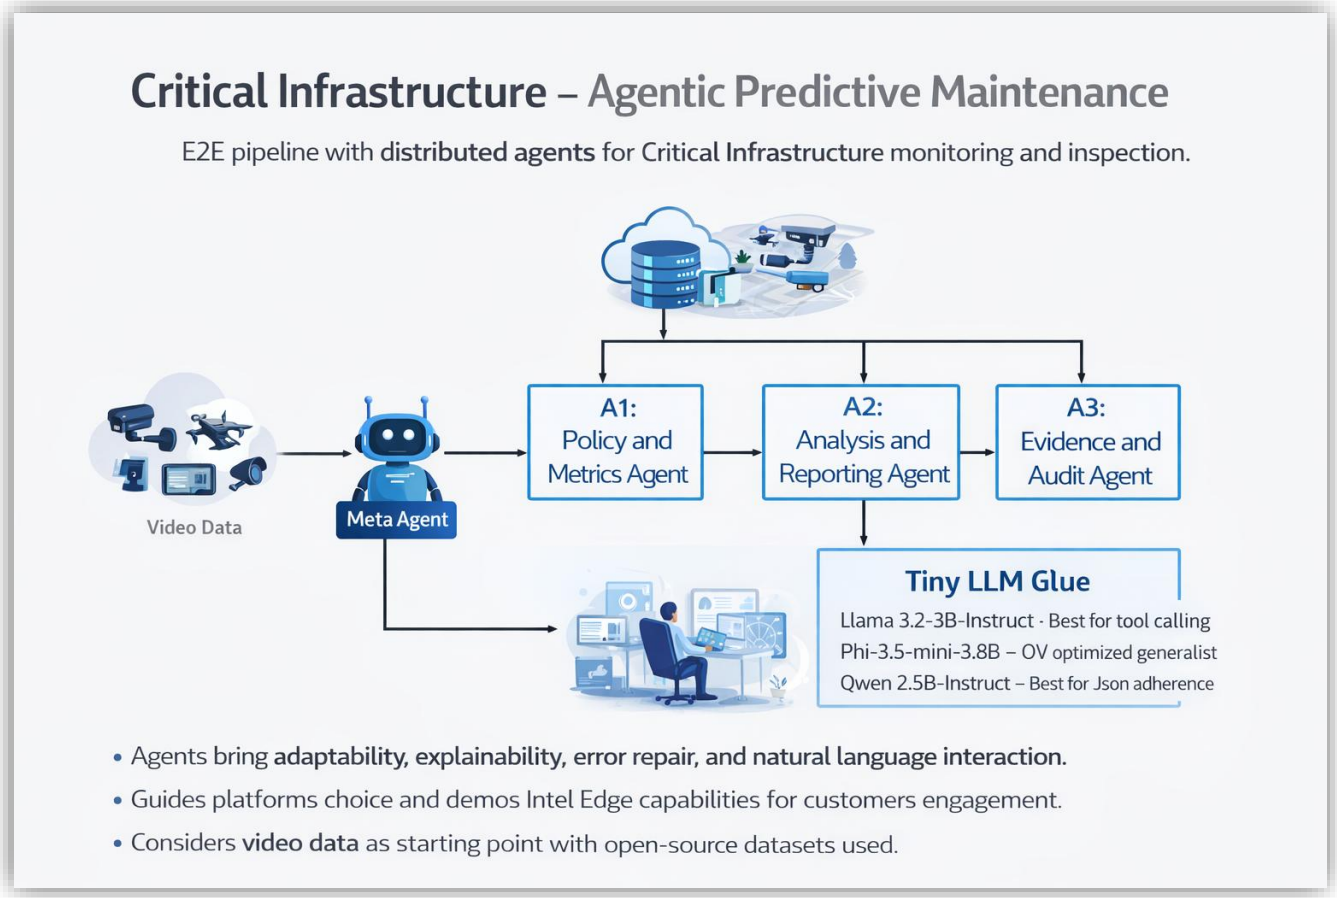

Key Features:

- **Performs real-time defect detection** from images and video at the edge while supporting industrial inspection scenarios
- Uses a **multi-agent reasoning pipeline** to analyze detections, assess severity, enforce policies, and generate structured, auditable outputs beyond simple AI inference

Supported Intel Platforms:

- Intel® CPUs with integrated graphics (12th Gen or later)

Explore details →

New

Multimodal Predictive Maintenance

Multi-modal Agentic pipeline powered by distributed AI agents for infrastructure monitoring and inspection leveraging data from sensors and cameras

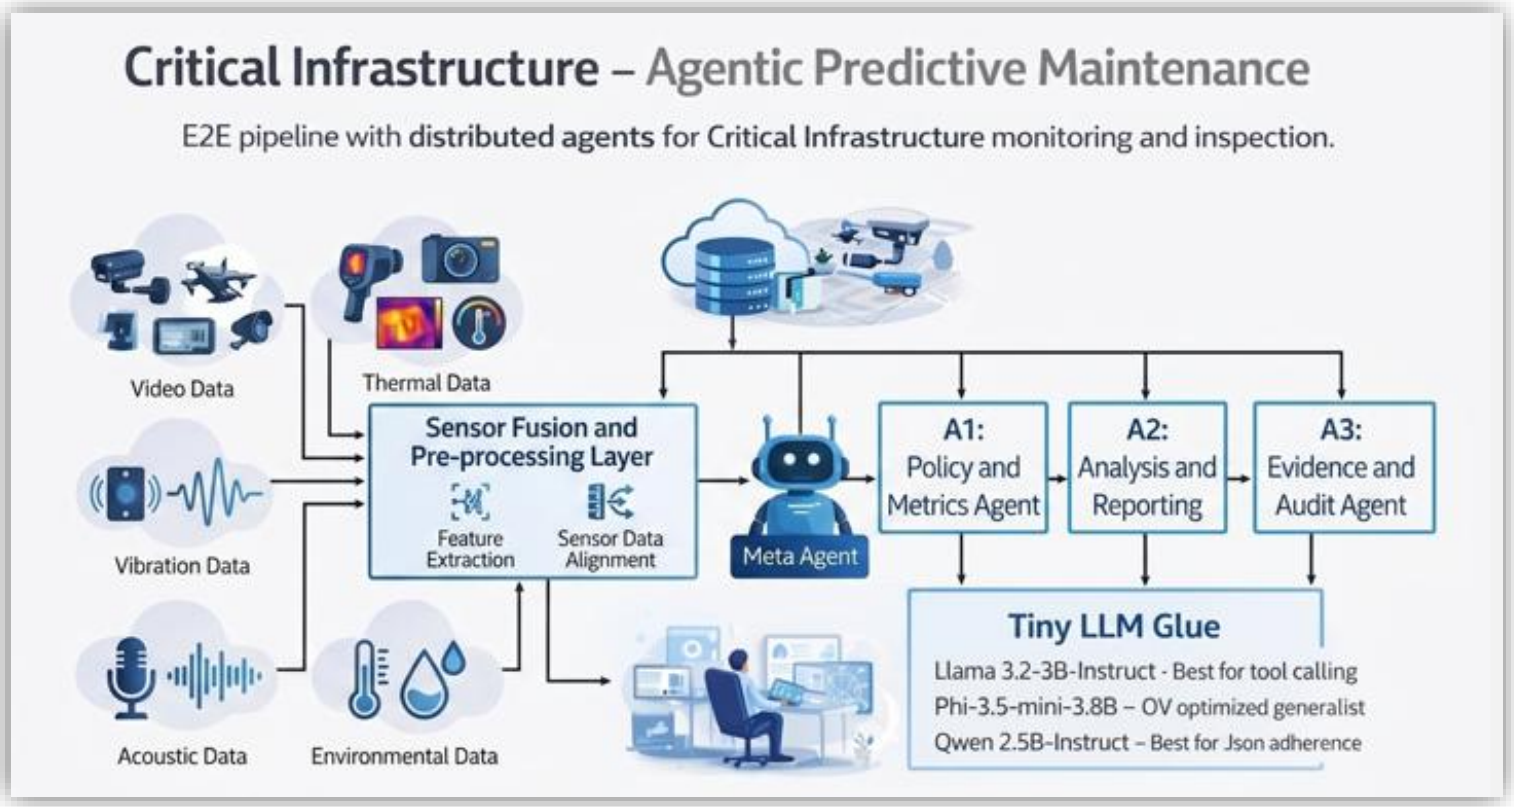

Key Features:

- Fuses edge camera vision with chemical sensors for a single, more reliable gas classification with fewer false positives
- Real-time multi-modal image and sensor data fusion that cuts false positives under ambiguous conditions and degrades gracefully when a camera or sensor fails

Supported Intel Platforms:

- Intel® CPUs with integrated graphics (12th Gen or later)

Explore details →

Smart Buildings Digital Twin Blueprint

Advanced pipeline that creates a high-fidelity virtual replica of smart buildings

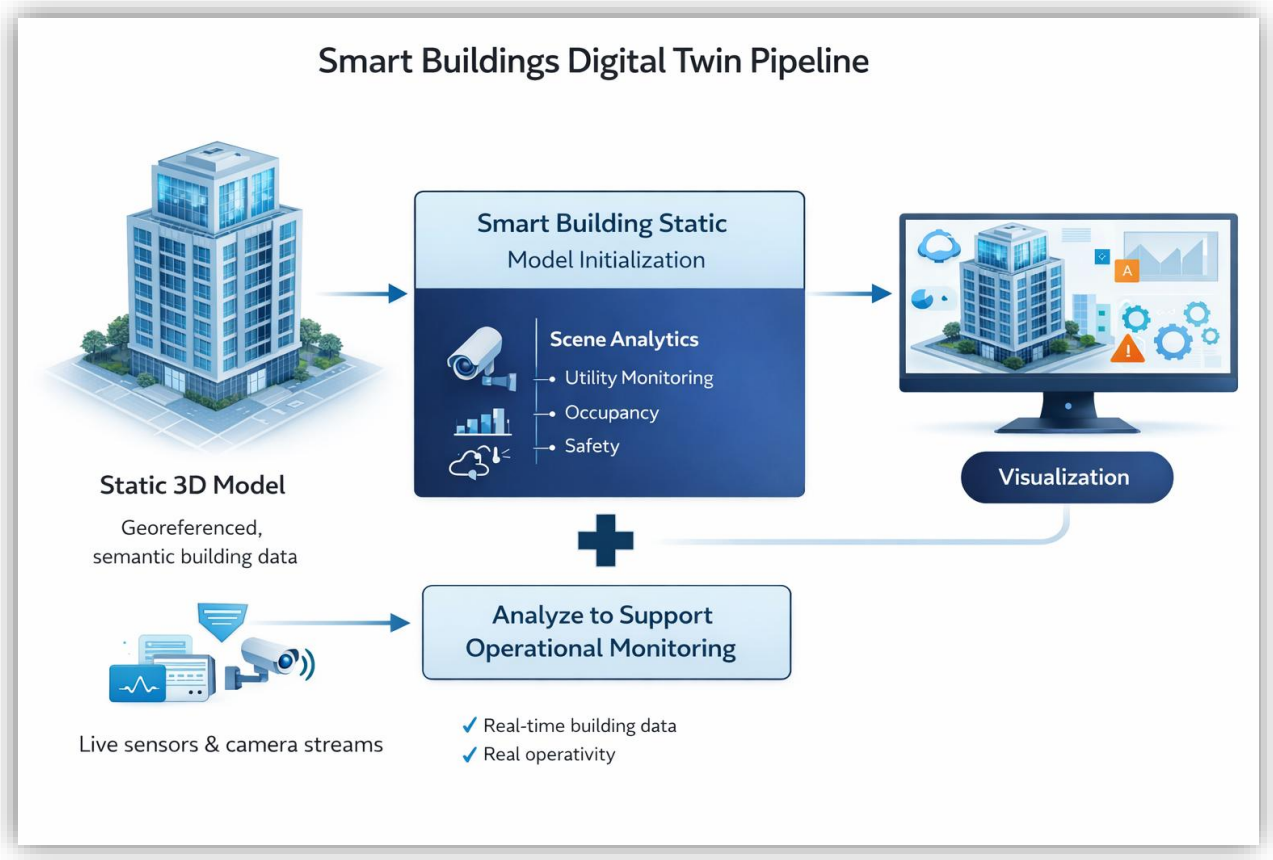

Key Features:

- Unifies multi-camera feeds into a single **4D digital twin** — a live metric 3D scene where every person, object, and door has precise position, velocity, and history
- Fuses edge vision AI with **environmental and identity sensor data** in real-world metric coordinates with SceneScape
- Runs **intelligent agents on scene data, not pixels** — powering access control, tailgating and object security, plus room occupancy, safety, and comfort

Supported Intel Platforms:

- 10th Gen or newer Intel® Core™ (i5 or higher); 2nd Gen or newer Intel® Xeon®

Explore details →

Enable ISV Customers

Metro AI Suite also gives Developers an easy, Optimized Path for Intel-based Edge AI

Metro AI Suite also includes materials for Software Developers

- SDK & documentation
- Sample apps & blueprints
- Tools to evaluate & measure performance on qualified systems

Intel offers support programs for ISVs and SIs, including:

- Metro AI Suite adoption
- AI architecture conversion
- Performance optimization

Open Source licensing

- Available for partners to use and market（Apache license）

Download Developer Guide

Next Steps

Metro AI Suite Validates, Qualifies, and Promotes System AI Capabilities

- Helps you quickly and easily **qualify your systems as AI ready**
- Enables your Software Developer customers to **quickly add Visual & GenAI to existing solutions**
- Demonstrate **maximized performance** of market leading edge AI silicon

Metro AI Suite Site

- intel.com/metroai
- Documentation
- Github link

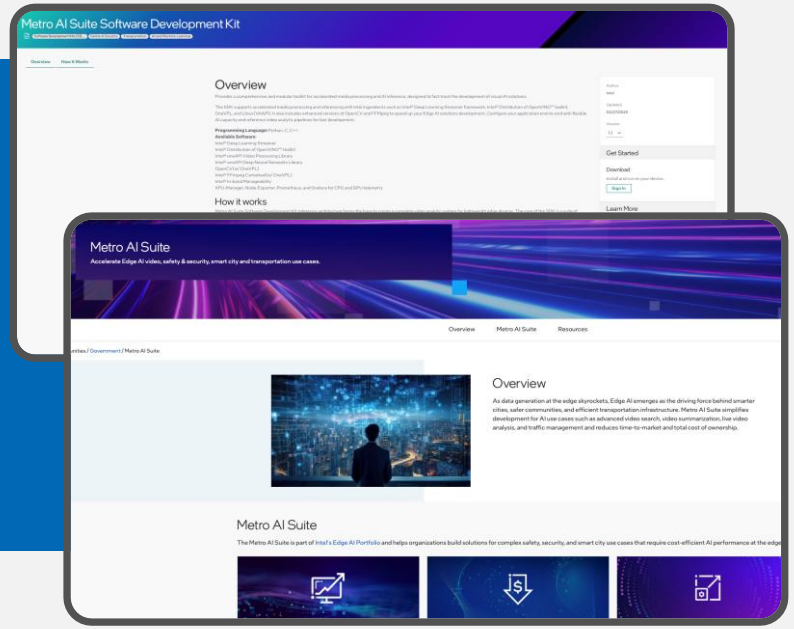

Next Steps:

- **Qualify your systems** with ESQ, then **Upload to Hardware Catalog**
- **Use Sample Apps & Rapid RI** to demonstrate AI system capabilities
- **Use ViPPET** to show Metro-enabled AI & media performance
- **Introduce your customers to Metro** SDK & programs