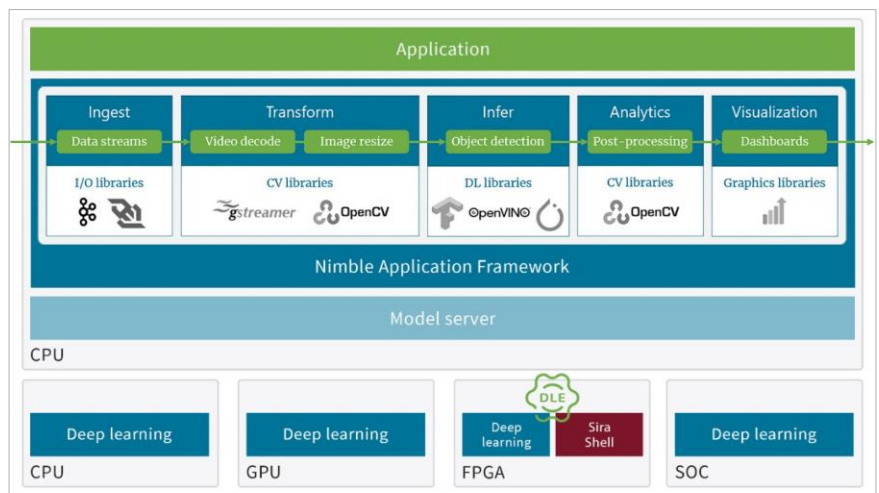# Real-time Streaming Analytics at the Edge

Empowering AI at the Edge: The Megh Open Analytics Platform, Powered by Intel® Technologies, Redefines Video Analytics for Outstanding Performance.

Contents

Executive Summary

Streaming data is experiencing exponential growth. Enterprises increasingly want to extract intelligence from this data, creating a growing demand for AI-based streaming analytics to process data at the edge.

The market for edge computing is growing rapidly. It is predicted that by 2025, over 50% of all data will be created at the edge.[1] This growth represents an enormous opportunity to provide solutions for streaming analytics from edge to cloud. There are, however, complex challenges in implementing edge analytics, including:

▪ Edge analytics applications must be capable of running on platforms with diverse hardware accelerators from multiple vendors, including CPUs, GPUs, FPGAs, and SOCs.

▪ The software architecture used for building edge analytics applications must be cloud native.

▪ Accurate edge analytics depend on continuously updated models.

1

Megh Computing developed Megh Open Analytics Platform to address these challenges. The platform is compatible with Intel® technologies and powered by the Nimble Application Framework, which enables easy implementation of analytics pipelines. It also supports a model server to accelerate AI algorithms using CPUs, GPUs, FPGAs, and SOCs. Megh Open Analytics Platform is based on Open Analytics (OA), an innovative approach to analytics emphasizing openness.

The platform has several key features:

▪ Edge-to-cloud deployments across hardware architectures with a single, unified deployment framework.

▪ Continuous training of pre-trained AI models with new data.

▪ Contextual analytics, in which existing cameras and sensors are used to consider the context of the data.

▪ Cross-platform support, which abstracts underlying hardware and frameworks, allowing applications to be transportable and pipelines to be distributed across hardware vendors.

Megh released Video Analytics Solution (VAS), an intelligent video analytics product line and the first to be powered by Megh Open Analytics Platform. VAS comes as a collection of market-ready solutions preconfigured for smart building, warehouse, retail, and other smart-place use cases. Megh's growing library of video analytics use cases is available, allowing any solution to be customized with support for physical security, worker safety, operational efficiency, inventory management, and traffic management use cases.

Megh's smart-place solutions consist of Megh's VAS Suite, VAS Portal, and VAS SDK products. Running at the edge, VAS Suite implements the Nimble Application Framework, enabling ingestion of data from multiple sources, data transformation, AI-based inferencing, and extracting event information.

Metadata from VAS Suite is forwarded to VAS Portal, which provides data visualization, notifications, queries, and reports. VAS SDK supports customization of the pipeline and visualization components.

Megh Open Analytics Platform enables users to:

▪ Eliminate false positives with advanced AI models optimized through real-time continuous training.

▪ Unlock new insights from existing cameras and sensors through contextual awareness created by ingesting diverse streams on the same platform.

▪ Reduce costs by leveraging cross-platform support across CPU, GPU, FPGA, and SOC architectures for the lowest TCO.

While Megh Open Analytics Platform excels as a foundation for intelligent video analytics, it has been designed to be extensible for analytics of any kind. Megh is developing Network Analytics Solution (NAS) to enable solutions that analyze network data and statistics.

Introduction: The Edge Computing Opportunity

Streaming data from billions of cameras, sensors, and other Internet of Things (IoT) devices is experiencing exponential growth. Information service provider IHS Markit estimated that, as of 2021, there were more than one billion surveillance cameras alone installed worldwide.[2] The rollout of private and public 5G networks is accelerating this growth.

These devices support applications in manufacturing, retail, healthcare, and other major industries, and are helping create "smart" buildings, factories, warehouses, and other data-driven places.

While common devices can generate gigabytes of data every day, a 2019 survey by Deloitte reported that only 18% of organizations were able to take advantage of unstructured data, like text, video, and audio.[3]

Enterprises increasingly want to extract intelligence from these data streams to gain actionable insights and create business value. Thus, there is a growing demand for AI-based streaming analytics to process data at the edge.

Edge computing refers to the range of computers, networks, and devices at or near the user. Meanwhile, edge analytics refers to the process of analyzing streaming data, often using AI in real time, before sending metadata to private or public data centers in the cloud for post processing. This enables rapid analysis and response.

Edge computing can be implemented at the device edge, an on-premises edge gateway, or at the telecommunications company (telco) edge. A typical solution is a hybrid approach implemented across different platforms from the edge to the cloud.

The market for edge computing is growing rapidly. Technology research and consulting firm Gartner predicts that by 2025, up to 75% of enterprise-generated data will be created at the edge.[4] According to Vantage Market Research, the global edge computing market was valued at $7.1 billion in 2021 and is projected to reach $49.6 billion by 2028, at a compound annual growth rate (CAGR) of 38.2% over the forecast period (2022-2028).[5] The edge AI software market is projected to grow from $0.8 billion in 2022 to $3.1 billion by 2027, at a CAGR of 29%, reports MarketsAndMarkets.[6]

This growth represents an enormous opportunity to provide solutions for streaming analytics from edge to cloud.

The Problem: Edge Computing Challenges

There are, however, complex challenges in implementing edge analytics. Three core problems relate to hardware compatibility, cloud platform compatibility, and AI model customization.

Diverse Hardware Platforms

Edge analytics applications must be capable of running on platforms with diverse hardware accelerators from multiple vendors, including CPUs, GPUs, FPGAs (field programmable gate arrays— flexible and fast accelerators), and SOCs (system on a chip—integrated circuits that integrate most or all components of a system). Applications need to be able to run across these architectures to provide deployment flexibility, low cost, and high performance. Moreover, application developers need a unified development framework to create solutions that scale easily.

Cloud-native Software Architecture

The software architecture used for building edge analytics applications must be cloud native, so that it provides a consistent development and deployment model that leverages the automated deployment, scaling, reliability, and failover capabilities available through containerization (packaging code and its dependencies so that it runs quickly and reliably across computing environments). This architecture facilitates software distribution across platforms and support for different native frameworks. For example, the frontend of an application could run on ARM processors at an edge router, with AI processing for analytics on an edge gateway using GPUs and post processing on CPU instances in the cloud.

Customized Deep Learning Models

Edge analytics application accuracy depends on the accuracy of the deep learning models that perform the analytics. In traditional approaches, a model is trained on a fixed dataset and deployed to run unattended. This can lead to inaccurate results over time as the environment changes. To deliver more accurate predictions and decisions, models should be updated continuously. They also need to be optimized to fit within the memory and computing constraints of the hardware platform (for example, deploying tiny models in edge devices, compared to larger models on edge servers or the telco edge).

Megh's Solution: Megh Open Analytics Platform

Megh Computing developed Megh Open Analytics Platform to address these challenges.

The platform is powered by the Nimble Application Framework, which enables easy implementation of analytics pipelines. Such pipelines consist of the following processing stages:

- Ingest any type of real-time stream, including audio, video, sensor, and network data.
- Transform the data, normalizing and storing it for subsequent processing.
- Infer information from the data using machine learning- and deep learning-based AI algorithms.
- Perform analytics to extract insights from the data.
- Visualize the data and configure alarms and notifications as desired.

Megh Open Analytics Platform also supports a model server to accelerate AI algorithms using CPUs, GPUs, FPGAs, and SOCs.
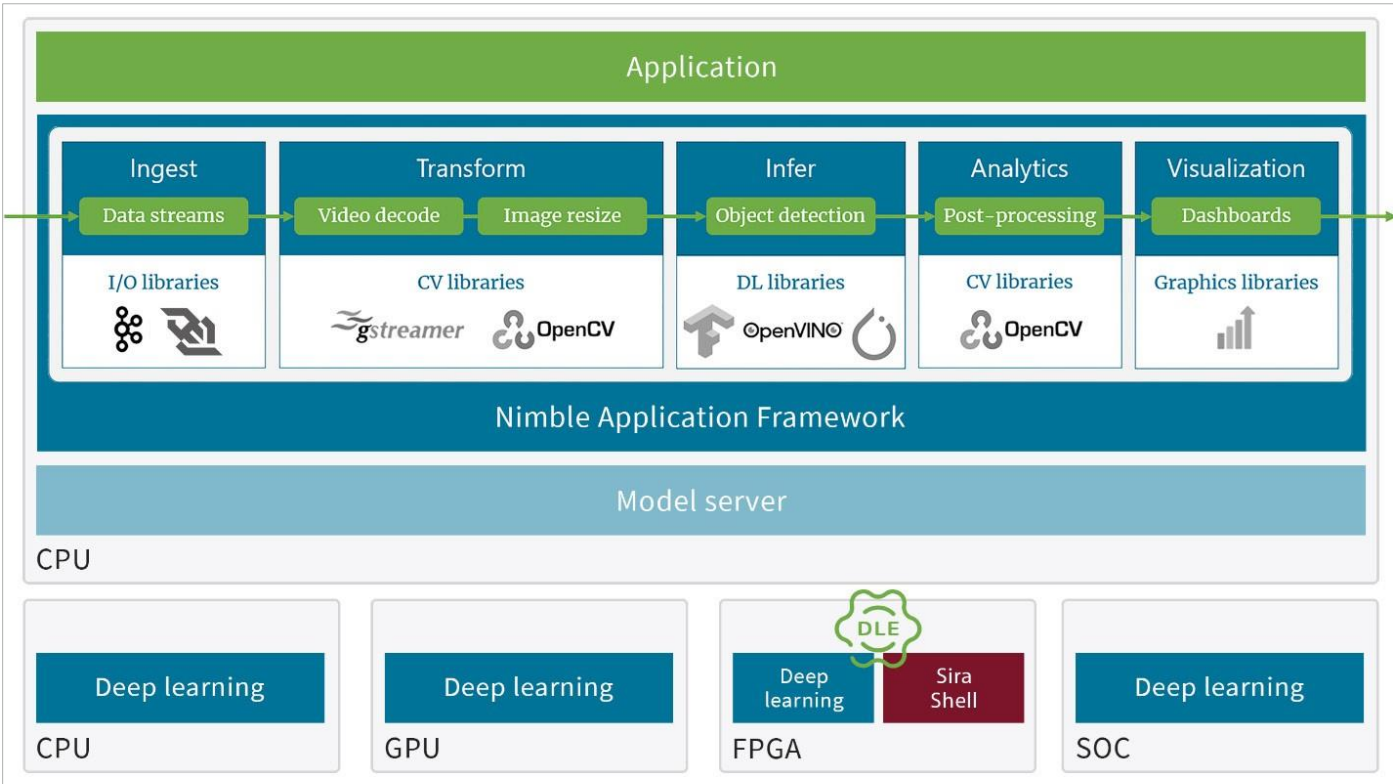
Open Analytics

Megh Open Analytics Platform is based on Open Analytics (OA), an innovative approach to analytics that gives business users, integrators, and developers the power to easily integrate real-time AI into their business operations, applications, and products to get actionable insights with operational reliability.

OA solutions are defined by the following three pillars.

Open Customization

Analytics solutions should support custom pipelines, so you can:

- Configure end-to-end pipelines.
- Swap or retrain/redeploy AI models for single- or multiple-stage inferencing.
- Configure parameters or deploy new libraries to implement business rules.
- Train AI models continuously.

Open Choice

Analytics solutions should work across compute contexts, supporting multiple:

- Hardware options (CPU, GPU, FPGA, and SOC).
- Vendors, through a single development framework.
- Deployment models (edge, on-premises, private and public cloud, and hybrid).

Open Integration

Analytics solutions should have cloud-native architecture, with easy integration with devices and open-source frameworks (such as FFMpeg, Gstreamer, TensorFlow, PyTorch, and OpenVINO), enabling sensor fusion at the edge. This includes:

- APIs to integrate with business systems (vendor management systems, business intelligence/ visualization, and more).

- APIs to integrate with business systems (vendor management systems, business intelligence/ visualization, and more).
- Multiple application and UI options (web, mobile, command line interface, and more).
- A software development kit (SDK) to embed functionality in any application or product.

Integral Intel Components

Megh's proprietary deep learning-based models can run on all Intel® hardware including Intel® Core™ processors, Intel® Xeon® Scalable processors, and Intel® NUC kits, and leverages the Intel® Distribution of OpenVINO™ toolkit. This technology optimizes Megh Computing's Open Analytics and VAS solutions, enabling developers to harness the full power of Intel® Silicon for real-time analytics. Megh builds custom deep learning models for object detection (people, vehicle, PPE) using open-source frameworks like TensorFlow and Pytorch. To maximize performance on Intel® CPUs, Megh takes these models through the Intel® Distribution of OpenVINO™ toolkit optimization flow to get the best performance and lowest power consumption on Intel® Silicon. This translates to the lowest total cost of ownership (TCO) for customers.

The Intel® Distribution of OpenVINO™

The Intel® Distribution of OpenVINO™ toolkit enables users to scale Megh's VAS to meet their deployment requirements and reduce the technical barriers to accelerating AI inferencing. Megh's deep learning extension (DLE) includes an FPGA plugin designed to integrate with the Intel® Distribution of OpenVINO™ toolkit. This extension broadens the toolkit's compatibility, making it suitable for deployment across a diverse range of Intel hardware platforms, including Intel® Core™ processors and Intel® Xeon® processors. This enables Megh Computing to develop real-time streaming analytics on less expensive edge devices and scale it up to more powerful devices when ready, without needing to refactor their solution.

With Megh's VAS, the barriers to unlocking real-time insights are drastically lowered for businesses of all sizes.

Intel® Xeon® Scalable Processors

Intel® Xeon® Scalable processors are specifically designed for demanding workloads and server applications. Intel® Xeon® processors offer enhanced performance, scalability, and reliability, making them well-suited for video analytics deployments that require higher processing power and efficient resource utilization. These processors support advanced features like Intel® Advanced Vector Extensions (Intel® AVX) and Intel® Turbo Boost Technology, which further accelerate computations and optimize performance.

Intel® Core® Processors

Megh Computing leverages Intel® Core™ processors, a product line of high-performance microprocessors, to achieve powerful computing at the edge for their video analytics applications. These processors are highly versatile as they offer their CPU's default computing capabilities or can leverage the CPU's integrated graphics processing unit (iGPU), for additional compute power. By leveraging both the CPU and iGPU to analyze video streams, Megh Computing can customize the required computing power of their solutions. They do this by adjusting the resolution and frame rate of camera streams, which enables scaling the solution to accommodate the analysis of more streams. This adaptability ensures that Megh Computing can deliver high-performance video analytics solutions that precisely meet their customers' needs.

Intel® NUC Kit

Megh Computing VAS is validated to run on Intel® NUC kits with 11th or 12th Gen Intel® Core™ processors. Intel® NUC kits offer an energy-efficient design and compact form factor, making this hardware optimal for deployment in edge devices.

This hardware is also available in various configurations, allowing users to select the processing power and memory capacity that suits their specific requirements. This scalability holds immense value for Megh's VAS applications, as it allows customers to adjust the hardware to meet the demands of their analytics workloads. In summary, Intel® NUC plays a key role in providing customers with the optimal balance of performance, power, and price to run and scale Megh VAS at the edge.

Megh VAS: Validated to Offer Scalable, High Performance on Intel® NUC Kits

Megh Computing and Intel worked together to validate the performance results of Megh Video Analytics Solution (VAS) running on Intel® NUC kits with 11th or 12th Gen Intel® Core™ processors. Through extensive testing, which included running Megh VAS on CPUs or on the integrated GPUs as accelerators, these results demonstrated the impressive capabilities of leveraging Intel® NUC kits to scale Megh VAS solutions with sustained performance. The results of this validation process are as follows:

The total number of streams supported by a system will depend on the camera's resolution (e.g., FHD 1080p or HD 720p) and frame rate in terms of Frames Per Second (e.g., 30 FPS or lower) of the video stream. Meanwhile, the required resolution and frame rate for a solution will depend on the use cases deployed. For instance, physical security use cases like intrusion detection may only require 720p resolution at 5 FPS, while license plate recognition use cases may require higher resolutions and frame rates. By simply managing both the resolution and frame rate in the processing pipeline, Megh's VAS is able to deliver very high

throughput on the Intel® NUC kits, while creating capacity for more streams to be analyzed. This empowers customers to leverage 24-hour, concurrent multi-stream video analytics that utilize the optimum amount of CPU or integrated GPU without compromising system accuracy.
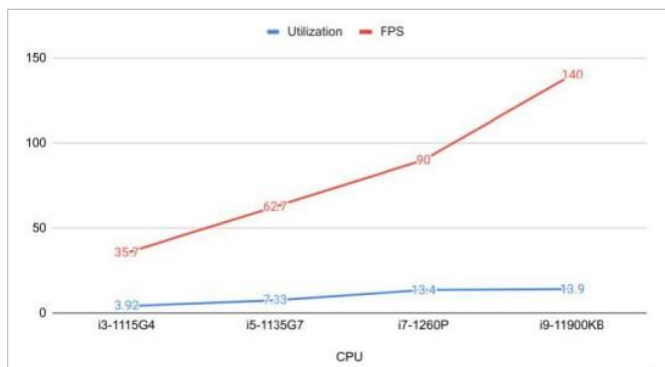
On average, with the integration of Intel® NUC kits, Megh Computing VAS users can analyze:

13 Video Streams with Intel® Core™ Processors.

22 Video Streams with the Integrated GPU.

35 Video Streams with the Intel® Core™ Processors and the iGPU.[7]

4.3 MAX FPS on CPU

When using only CPUs to run the video analytics pipeline, the maximum FPS you can achieve increases with more CPU cores. For instance, Megh VAS starts at 35 FPS with an Intel® Core™ i3 processor and goes up to 140/150 FPS with an Intel® Core™ i9 processor.[7] Even as the FPS goes up, the CPU is not heavily used, leaving available computing power for other tasks.

4.4 MAX FPS on iGPU
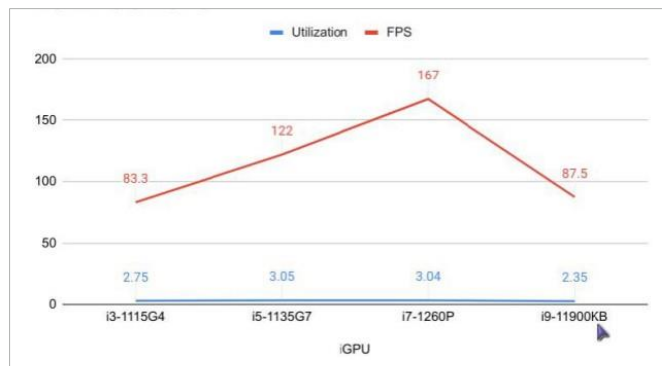
Observation: The lower FPS of the iGPIU on the Intel® Core™ i9-11900KB processor compared to the Intel® Core™ i5 and Intel® Core™ i7 processors was caused by the lower number of execution units in the CPU package (of the Intel® Core™ i9 processor).

The maximum FPS achievable when using the integrated graphics (iGPU) on Intel® Core™ processors for inference, along with CPU utilization for video stream decoding and subsequent analytics, varies depending on the type of processor. On an Intel® Core™ i3 processor, it can perform 83 FPS while the Intel® Core™ i7 processor scales the Megh VAS solution to 167 FPS. However, with the Intel® Core™ i9 processor, the FPS lowers to 87 FPS because there are fewer iGPU execution units on this processor type.[7] Therefore, the Intel® Core™ i7 processor offers the best performance for the least cost, making this the optimal hardware for this solution.

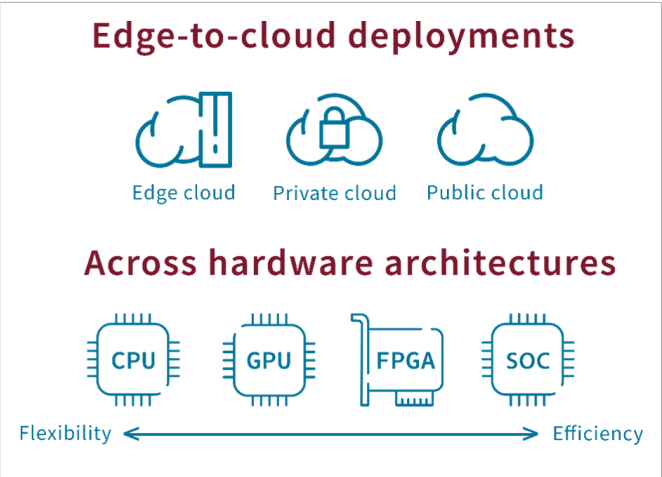View the full Megh VAS performance and validation report on Intel® NUC kits here.

Platform Features

Megh Open Analytics Platform meets the challenges of edge analytics, delivering best-in-class performance, thanks to several central platform features.

Universal Deployment
Megh Open Analytics Platform supports edge-to- cloud deployments across hardware architectures with a single, unified deployment framework.

The platform also supports a comprehensive deployment infrastructure to allow customers to update and monitor sites remotely.

Edge-to-cloud deployments

Edge cloud   Private cloud   Public cloud

Across hardware architectures

CPU   GPU   FPGA   SOC

Flexibility ←————————→ Efficiency
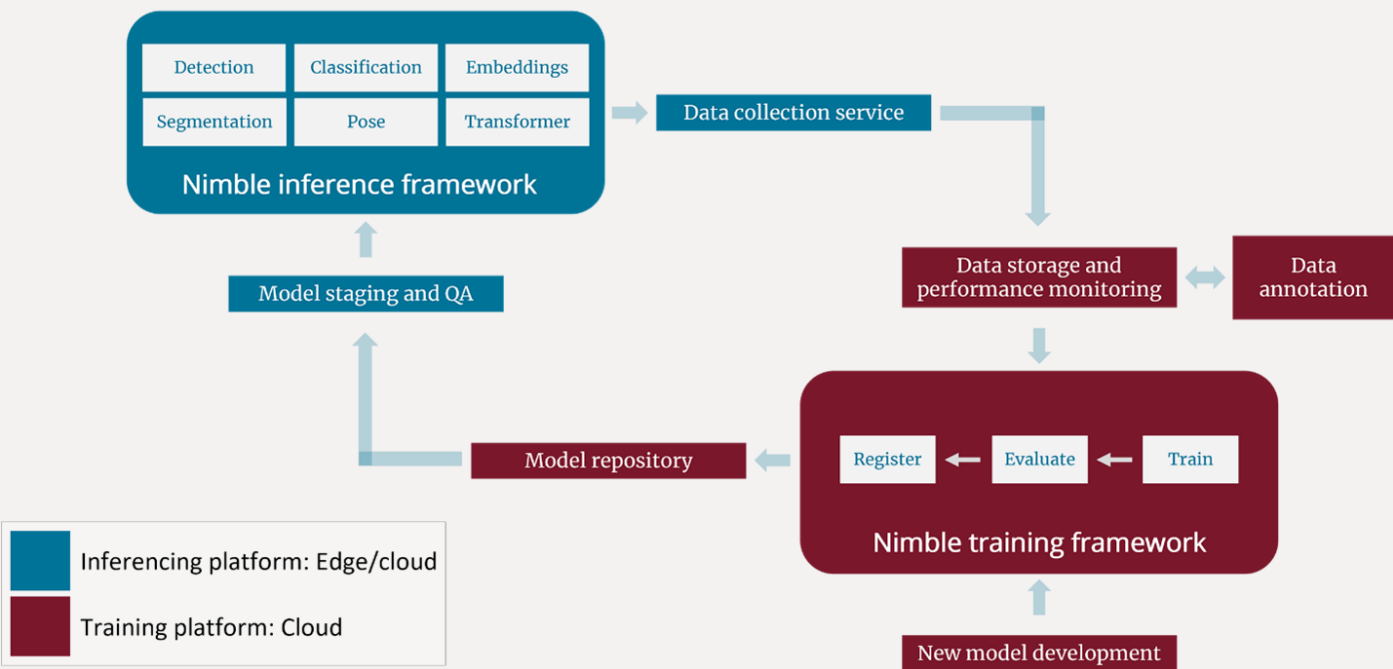
Continuous Training

Megh's platform supports its continuous training (CT) framework for constant fine tuning of pre- trained AI models with new data, allowing them to adapt and improve continually.

The result is more accurate predictions and decisions, especially as the environment in which models are deployed changes. The CT framework allows models to adapt quickly and dynamically, virtually eliminating false positives.

The framework comprises the following components:

1. Inferencing service: Nimble with on-premises inference engine processes video frames in real time.

2. Data collection service: Selected frames with lower confidence are forwarded to a data lake for storage, performance monitoring, and annotation.

3. Continuous training: Retrain existing or develop new models on-demand using the data lake and evaluate historical performance for potential   deployment.

4. Model staging: Stage new models for deployment and evaluate their performance on real-time data vs. existing models.

Megh Platform's continuous model optimization

Detection | Classification | Embeddings
Segmentation | Pose | Transformer
Nimble inference framework

Data collection service

Model staging and QA

Data storage and performance monitoring

Data annotation

Model repository

Register ← Evaluate ← Train

Nimble training framework

New model development
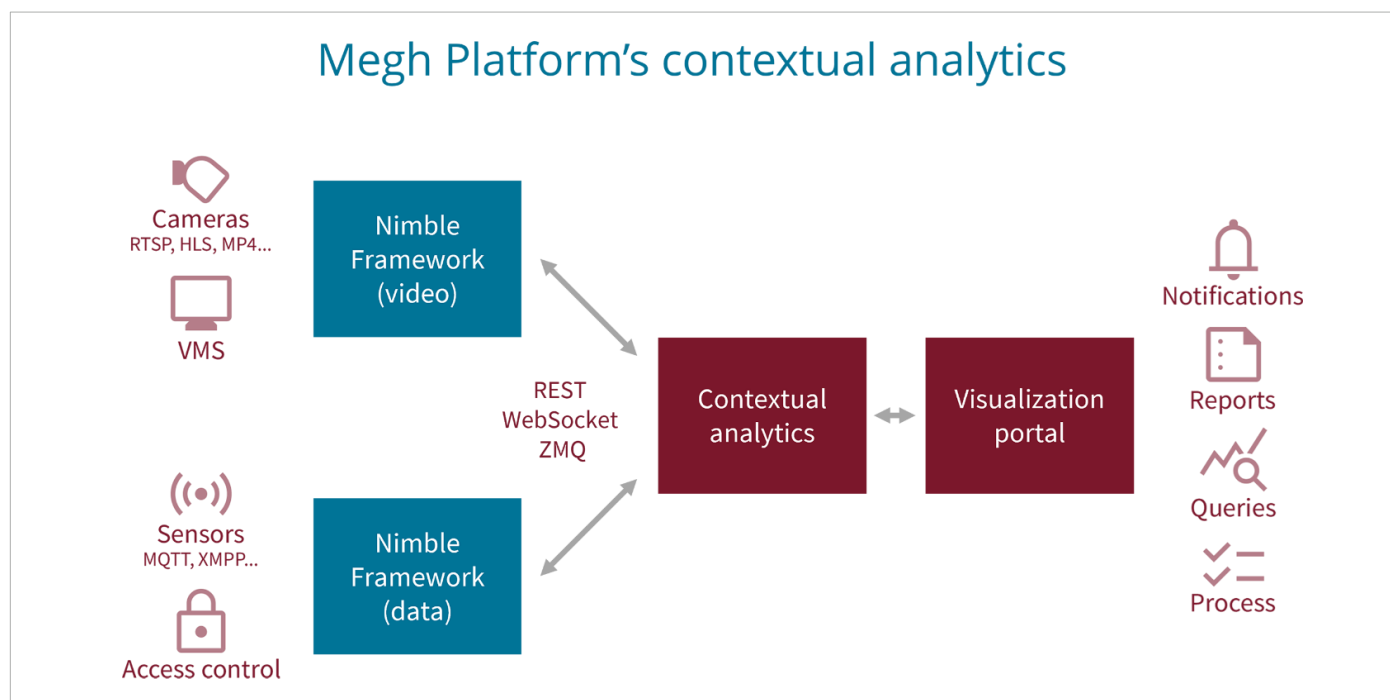
Inferencing platform: Edge/cloud
Training platform: Cloud

This approach offers several important benefits:

- Fast modeling: Optimized MLOps pipeline implemented with data collected on site and trained in the cloud, resulting in new models in four to six weeks.
- Accuracy: Training data collected from inferencing service dramatically improves model accuracy and enables continuous monitoring to ensure robust model performance.
- Efficiency: Models implemented at the edge for inferencing use a range of hardware architectures (CPUs, GPUs, FPGAs, and SOCs).

Contextual Analytics

The platform also supports Megh's contextual analytics (CA) framework, which considers the context in which data is generated, collected, and analyzed to provide more complete and accurate understandings. The framework unlocks new insights from existing cameras and sensors,   providing a new level of intelligence and efficiency. A growing framework, the rule-based engine for CA will be enhanced with generative AI to provide a conversational interface for situational awareness using sensor fusion.

Megh Platform's contextual analytics

Cameras RTSP, HLS, MP4…
VMS
Sensors MQTT, XMPP…
Access control
Nimble Framework (video)
Nimble Framework (data)
REST WebSocket ZMQ
Contextual analytics
Visualization portal
Notifications
Reports
Queries
Process

VAS Suite supports ingesting streaming data, normalizing it for storage in real-time databases, and performing analytics on individual streams to determine events using the Nimble Application   Framework. VAS Portal supports an engine implementing the contextual rules for alerts, which can be based on a combination of events across streams. The engine supports a combination of machine learning and statistical methods, such as regression analysis and decision trees, to define contextual rules. This is followed by powerful, query- based data visualization techniques.

The result is a more nuanced understanding of the data, which can help organizations make more informed decisions and take more effective action.

Cross-platform Support

Megh Open Analytics Platform supports Megh's cross-platform support (CPS) framework, which hides the complexities of underlying hardware and frameworks from different vendors. This allows streaming applications to be transportable and pipelines to be distributed across systems, so application developers can easily choose the right hardware.

Megh continually enhances the CPS framework through ongoing development   of the Nimble Application Framework, providing comprehensive support for development of edge analytics applications.

The CPS framework has the following key features.

REST and WebSocket APIs

Nimble's REST and WebSocket APIs provide a well- defined abstraction for managing high-level analytics pipelines, while also allowing fine-grained control over how those pipelines are executed across a range of   available hardware resources.

Secure Data Transmission (Encryption + Authentication)

Nimble supports distributed edge computing out of the box, thus encryption and authentication is required for securely transmitting data from edge to cloud.

Resource and Execution Management

Nimble handles the complex task of managing and distributing execution work efficiently across hardware resources. At its core it is an advanced thread-pooled task scheduling engine that takes advantage of a pipeline element's metadata to know when the element can be processed in parallel.

Software Development Kit

Nimble empowers developers to create custom pipelines by providing a comprehensive SDK for incorporating new models and accelerating them across devices. The SDK also provides a flexible pipeline element interface that gives developers the ability to create simple and complex analytics. These models and analytics can be written in a variety of programming languages.

Hardware Acceleration Support

Nimble also includes an API for hardware accelerators. With this API, developers can easily integrate platform hardware accelerators with Nimble and expose them as a pipeline element. The plugin framework is built from the ground up to maximize performance.
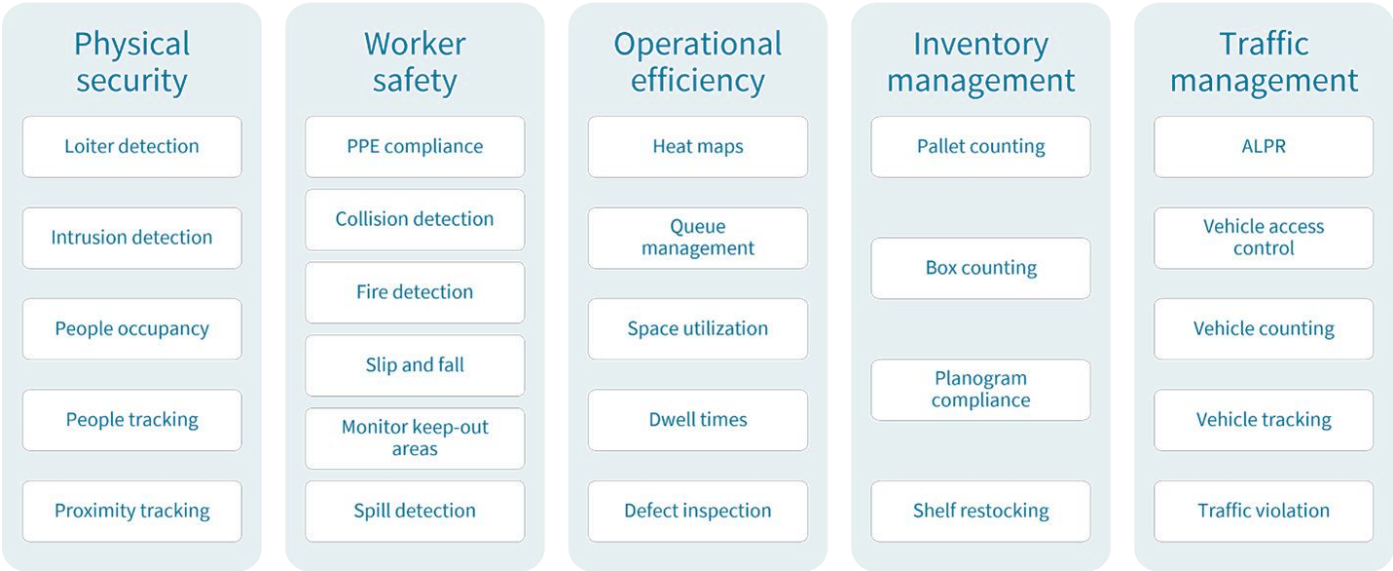
Video Analytics Solution

To enable business users, integrators, and developers to take advantage of the growing opportunities in edge analytics, Megh released Video Analytics Solution (VAS), an intelligent video analytics product line and the first powered by Megh Open Analytics Platform.

VAS comes as a collection of market-ready solutions preconfigured for smart building, warehouse, retail, and other smart-place use cases to reduce risk and improve operational efficiencies.

| Smart BUILDINGS	Smart WAREHOUSES	Smart RETAIL	Smart CITIES	Smart FACTORIES
❖ Loiter detection ❖ Intrusion detection ❖ People counting	❖ Inventory tracking ❖ PPE compliance ❖ Collision detection	❖ Planogram compliance ❖ Inventory restocking ❖ Heat maps	❖ People management ❖ Traffic management ❖ Anomaly detection	❖ Quality control ❖ Safety monitoring ❖ Process optimization

VAS use-case library

Physical security	Worker safety	Operational efficiency	Inventory management	Traffic management
Loiter detection	PPE compliance	Heat maps	Pallet counting	ALPR
Intrusion detection	Collision detection	Queue management	Box counting	Vehicle access control
People occupancy	Fire detection	Space utilization		Vehicle counting
People tracking	Slip and fall	Dwell times	Planogram compliance	Vehicle tracking
Proximity tracking	Monitor keep-out areas	Defect inspection	Shelf restocking	Traffic violation
	Spill detection			

Megh's growing library of video analytics use cases is available, allowing any solution to be customized with support for physical security, worker safety, operational efficiency, inventory management, and traffic management use cases.

Solution Architecture

Megh's smart-place solutions consist of Megh's VAS Suite, VAS Portal, and VAS SDK products. Running at the edge, VAS Suite implements the Nimble Application Framework, enabling ingestion of data from multiple sources, data transformation, AI-based inferencing, and event information extraction from the processed data. Metadata from VAS Suite is forwarded to VAS Portal, which provides data visualization, notifications, queries, and reports. VAS SDK is available for the customization of the analytic pipeline and visualization components.
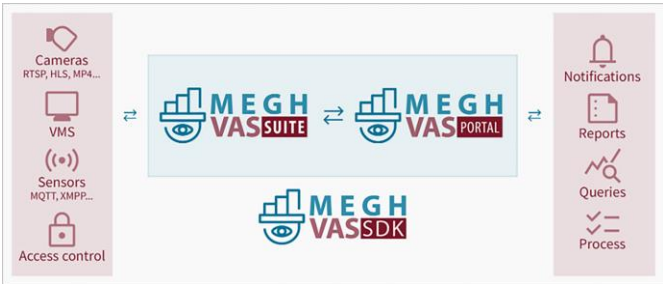
Typical Deployment

A common deployment configuration locates VAS Suite—implementing the video analytics use cases— at the edge and VAS Portal—to manage multiple VAS Suite instances—in the cloud.
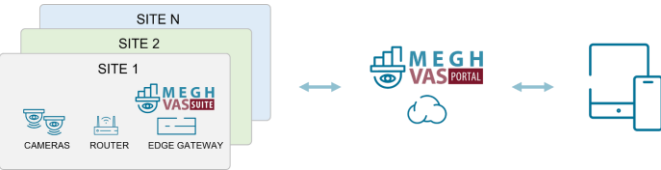
Alternatively, VAS Suite can also be deployed in the cloud.
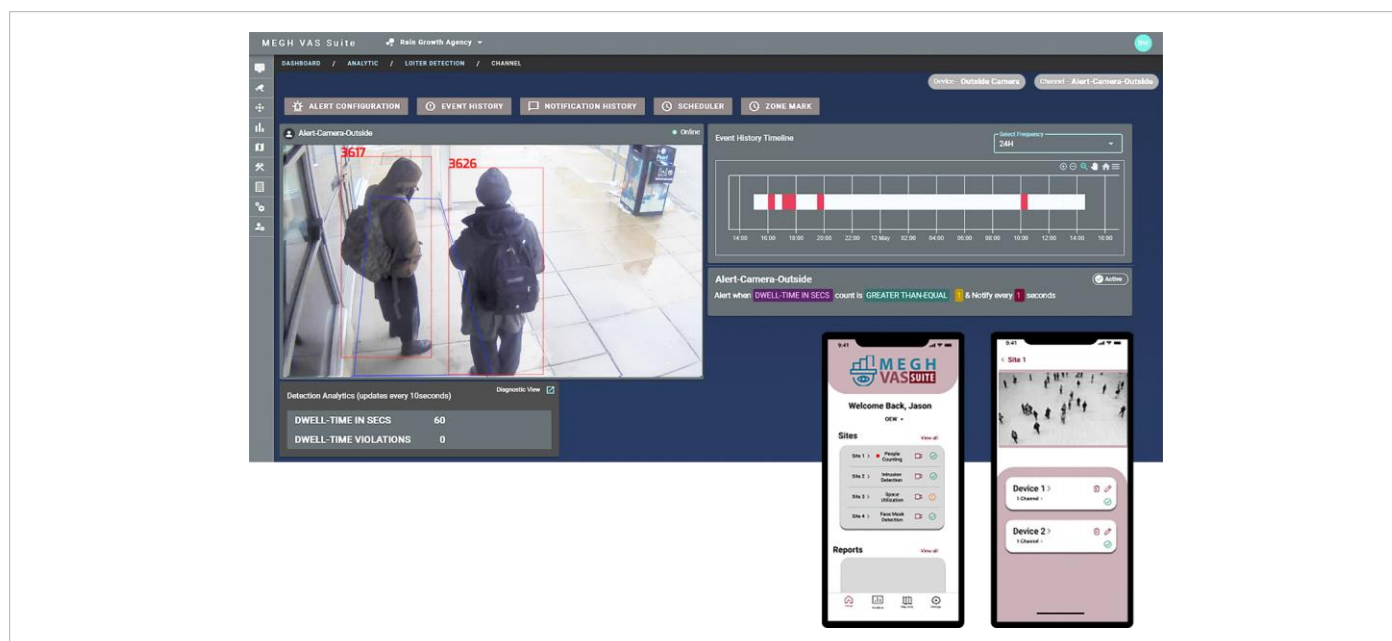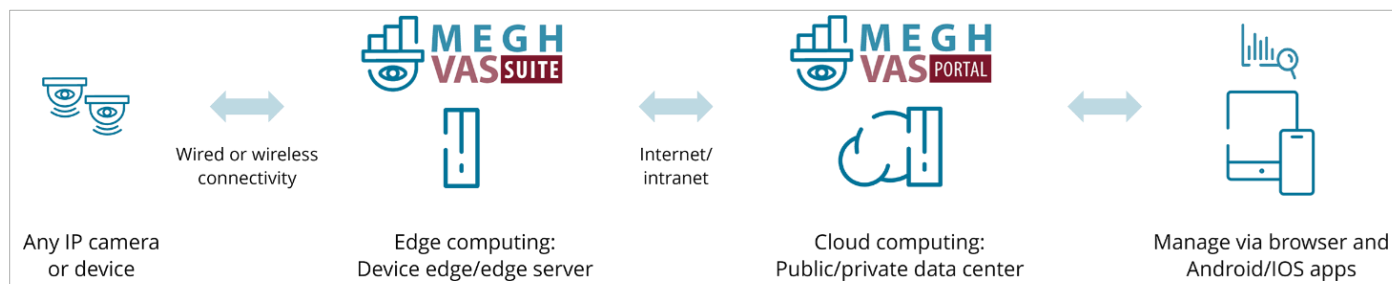
End-to-end Solution

VAS provides an end-to-end solution, encompassing devices, connectivity, analytics, information management, and system configuration.

- Any IP camera or device can become a smart device.

- Devices can stream using wired or wireless connectivity, including 5G networks.

- VAS Suite runs on the device edge, an edge server, or a hybrid configuration.

11

- Metadata is forwarded to VAS Portal in the cloud, running in public or private data centers.

- VAS Portal provides an access-anywhere dashboard for configuring cameras, alerts, notifications, queries, and reports.

Notifications can be sent in real-time to native Android and IOS mobile applications.

The complete end-to-end solution can be implemented from edge-to-cloud to support use cases for smart place verticals.

| Any IP camera or device | Wired or wireless connectivity | Edge computing: Device edge/edge server | Internet/ intranet | Cloud computing: Public/private data center | Manage via browser and Android/IOS apps |

Benefits

Megh Open Analytics Platform provides several critical edge-analytics benefits. It enables users to:

- Eliminate false positives with advanced AI models optimized through real-time continuous training.

- Unlock new insights from existing cameras and sensors through contextual awareness created by ingesting diverse streams on the same platform.

- Reduce costs by leveraging cross-platform support across CPU, GPU, FPGA, and SOC architectures for the lowest TCO (total cost of ownership).

From Video to Network Analytics and Beyond

While Megh Open Analytics Platform excels as a foundation for intelligent video analytics, it has been designed to be extensible for analytics of any kind. Megh is developing Network Analytics Solution (NAS) to enable solutions that analyze network data and statistics to identify trends and patterns. These solutions will support use cases like congestion control and intrusion detection to manage network performance and security. This will allow the same platform to address both physical and cyber security as the industry drives convergence of operational technology (OT) and IT in the enterprise.

Megh Computing & Intel in Action

Fixed-to-Mobile, AI-powered Video Analytics for Buildings

Pedro Pereira, **Intel insight.tech**
November 16, 2022

Smart buildings deliver continuous streams of data from sensors, cameras, and other devices to optimize operations and respond to user needs. With AI-powered video analytics, building operators can easily process this data. But these building solutions are usually within permanent structures supporting fixed surveillance systems. Many locations, including construction sites, temporary work sites, parks, event venues, and other in-field settings, need mobile surveillance units because they are typically not equipped with video cameras or analytics.

One company facing this situation turned to Megh Computing, a video analytics solution provider, about a Mobile Surveillance Unit (MSU). It needed an advanced, cost-optimized, AI-based video analytics solution to monitor 5,000 mobile sites, all on solar power.

Existing MSUs typically use expensive "smart cameras" with analytics that detect intrusion. Since the AI models used in these devices are not very robust, they generate many false positives, which increase operational costs in monitoring the alarms and the bandwidth required for the backhaul.

Megh Computing introduced a new MSU, with an Intel® Core™ i3 NUC processing streams from four standard cameras on the unit. The system uses advanced deep learning models based on an Intel® Distribution of OpenVINO™ Toolkit-supported inference engine for detection of people and objects.

Prabhat K. Gupta, CEO of Megh Computing, says, "Since these models deliver high operational reliability, Megh is able to virtually eliminate incidents of false positives, reducing operational costs for monitoring."

The MSU's video analytics pipeline is highly optimized and controlled by a motion filter that reduces overall power consumption for the system to less than the required 15W. This system also generates savings from lower communication and camera costs, resulting in compelling ROI.

"You want to be able to monitor any situation or crowd to see what's happening. With an MSU, you can quickly roll the thing up, put it up, and start monitoring in conjunction with security staff," Gupta says.

Customized Smart Building Solutions

Megh Computing leveraged its expertise with its VAS (Video Analytics Solution) for Smart Buildings to create the mobile option for its customers. VAS is a highly customizable solution that can target different hardware architectures and can be deployed from the edge to the cloud. VAS processes data from both cameras and sensors. It has been deployed in a variety of environments, including smart cities, smart warehouses, and retail locations.

VAS applies AI deep learning to accurately detect people and objects, and also anomalies in human behavior, to reduce physical risks. This was an important aspect to bring to its MSU solution.

"Obviously, there's an explosion of data at the edge, and most of the data is video," says Gupta. "And, except for data used to respond to security incidents, most of it is never analyzed."

Companies that invest in costly video systems want to be able to analyze streaming data and gain business insights from it. For instance, how many people use a company's spaces at any given time,

and for what purpose? The answers can lead to operational improvements, such as better traffic flow and smarter security controls. Think of a retail kiosk. By monitoring the space, Megh's solution can determine how long people wait in line, so the retailer can add another kiosk to avoid lost sales.

"We can also monitor space usage," says Gupta. "We can see how people are spending time in work areas throughout the day. That helps with facility planning, for example."

Reduce False Alarms with AI-based Security Solutions

Besides offering customization and mobility, Megh Computing addresses a common problem with video analytics: the high rate of false alarms. A security firm turned to Megh when it was experiencing this problem.

After piloting Megh's solutions side by side with their existing platform, the company made the switch. "The operator said we were basically able to eliminate all of their false positives," Gupta explains.

Megh Computing minimizes false alarms through the use of its advanced AI deep learning models and its ability to continuously update these models for improved accuracy, using its continuous training framework. "This is a unique capability that assures that when the system issues an alert, the person who receives it can feel confident it's real and take appropriate action to prevent a security incident," says Duncan Moss, Principal Engineer at Megh Computing.

For instance, a car dealership uses the platform for after-hours monitoring. If someone acts suspiciously on the lot, intelligent cameras can guess the person's intent by tracking their movements and alerting a security guard. An individual who moves from car to car, stopping for a few seconds by each vehicle, may plan to break in and steal a car. "That's an example of how we are using a technology for behavioral analysis and prevention," Moss says.

End-to-end advanced Security Solutions

Megh Computing not only looks at situational monitoring of spaces but aims to address cybersecurity in its solutions. "People who are trying to break in can break in physically or electronically. If you want to provide complete security for an enterprise, you've got to look at both aspects," Megh CEO Gupta says.

For example, VAS addresses the cyber aspect by monitoring network packets for congestion and signs of distributed denial of service (DDoS) activity. Thus, Megh provides a single platform for both physical security and cyber security threats.

To deliver these capabilities, Megh Computing partners with Intel for its advanced AI hardware and software platforms. Megh's solutions are certified as Intel Market Ready.

Going forward, Megh plans to continue leveraging the Intel relationship, not just for technology but also to widen its market reach. Having gained traction in smart buildings, Gupta says the company now targets retail and smart warehouses, too.

As it does, Megh continues to look for opportunities to deploy MSUs. After all, security and safety shouldn't be confined to buildings; it's needed wherever there is risk.

intel.

Source

[1]Goasduff, Laurence. 12 Data and Analytics Trends to Keep on Your Radar. Gartner, April 2022. 12 Data and Analytics Trends for Times of Uncertainty (gartner.com)

[2]Cosgrove, Elly. One billion surveillance cameras will be watching around the world in 2021, a new study says. CNBC, December 2019. One billion surveillance cameras will be watching globally in 2021 (cnbc.com)

[3]Harbert, Tam. Tapping the power of unstructured data. MIT: February 2021. Tapping the power of unstructured data | MIT Sloan

[4]Equinix is driving the digital future. Equinix. https://www.equinix.com/resources/analyst-reports/gartner-distributed-enterprise-predictions-2022

[5]Vantage Market Research. Growth: Edge Computing Market Size & Share to Surpass $49.6 Billion by 2028. Globe Newswire, December 2022. Growth: Edge Computing Market Size & Share to Surpass $ (globenewswire.com)

[6]Edge AI Software Market by Component, Data Source, Organization Size, Vertical and Region – Global Forecast to 2027. MarketsAndMarkets. Edge AI Software Market Size, Share and Global Market Forecast to 2027 | MarketsandMarkets

[7]Megh Computing VAS Performance and Validation Report for Integration on Intel NUC Kit – Report. Intel, September 2023. Megh Computing VAS Performance and Validation Report for Integration on Intel NUC Kit - Report

Notices & Disclaimers