Lenovo Validated Design for Smart Cities with VSaaS.ai

Provides a technical overview of Lenovo ThinkEdge servers

Highlights VSaaS.ai Smart Cities Use Cases

Explains the benefits of Edge AI solutions for Smart Cities

Provides solution sizing models and validation results

Vanita Meyer
Chris Arenas
Francisco Soto
Rodolfo Castillo
Andres Ulloa

LENOVO PRESS

Table Of Contents

1. Introduction

What makes cities smart?

Smart cities leverage data and digital technology to enhance the quality of life by enabling real-time monitoring and rapid, cost-effective responses to changing demand patterns. Smart technologies transform infrastructure by lowering the cost of data collection, allowing city governments, businesses, and residents to optimize existing systems. These technologies also engage the public in shaping demand, encouraging behavior such as using transit during off-peak hours or conserving resources. The outcome is a city that is not only more pleasant but also more conducive to business productivity.

Smart cities are increasingly adopting AI and edge computing technologies to enhance urban management, security, and safety to improve the quality of life for residents. Key trends include:

- **AI and Machine Learning Integration**: Advanced analytics capabilities such as real-time object recognition, behavior analysis, and predictive analytics are being leveraged to enhance threat detection and automated response.
- **Edge Computing**: Local data processing at the edge reduces latency and bandwidth usage, enabling faster decision-making in critical situations.
- **IoT Integration**: Video analytics is being integrated with IoT devices like smart streetlights and environmental sensors to create a comprehensive urban monitoring network.
- **Smart Traffic Management**: Video analytics are used to monitor traffic flow, detect accidents, and optimize traffic signals, improving transportation efficiency and road safety.
- **Public Health Monitoring**: Video analytics assists in public health initiatives by monitoring crowd density, mask usage, and temperature screening to prevent the spread of diseases.
- **Facial Recognition Technology**: Facial recognition is being implemented for enhanced security and efficient identification.
- **Sustainability Initiatives**: Computer vision systems help monitor environmental conditions, detect pollution, and ensure regulatory compliance, supporting sustainability efforts.
- **Proactive Maintenance and Infrastructure Monitoring**: AI-equipped video analytics monitors the condition of critical infrastructure, enabling proactive maintenance and ensuring the longevity and safety of urban assets.

These trends showcase the pivotal role of AI, video analytics, and edge computing in the evolution of smart cities, driving advancements in safety, security, efficiency, and overall urban living conditions.

This document offers planning and design considerations, along with best practices for implementing smart cities solutions using Lenovo ThinkEdge products and VSaaS.ai solutions.

Intended Audience

The intended audience for this document includes IT professionals, technical architects, sales engineers, channel partners, and consultants. It aims to assist in the planning, designing, and implementation of Lenovo and VSaaS.ai smart cities solutions.

2. Business challenges and business value

2.1 Business problem

According to the World Bank, as of today, 4.4 billion people, representing 56% of the global population, live in urban areas. By 2045, this number is expected to rise to 6 billion, adding nearly 2 billion more urban residents. With cities generating over 80% of global GDP, urbanization holds significant potential to drive sustainable growth. If managed effectively, urbanization can enhance productivity, foster technological innovation, and create an environment where new ideas can thrive, contributing to the overall economic and social well-being of populations worldwide.

As population grows, cities face numerous challenges that affect the quality of life for their residents. These problems include traffic congestion, high crime rates, poor air quality, inefficient public services, and environmental degradation. Lenovo ThinkEdge and VSaaS.ai partner to provide solutions that address these issues through various smart applications.

Source: World Bank Group

2.2 Business value

Cities generate vast amounts of data, including input from numerous existing cameras. Extracting insights from this data enables cities to respond quickly, allocate resources wisely, and plan effectively. Real-time information empowers them to make better decisions and actively shape the city's performance. As cities become smarter, leveraging existing infrastructure, data, and AI technologies to achieve safer, more efficient, responsive, and sustainable environments for the residents.

Figure 1. Smart city applications can improve some key qualities of life.

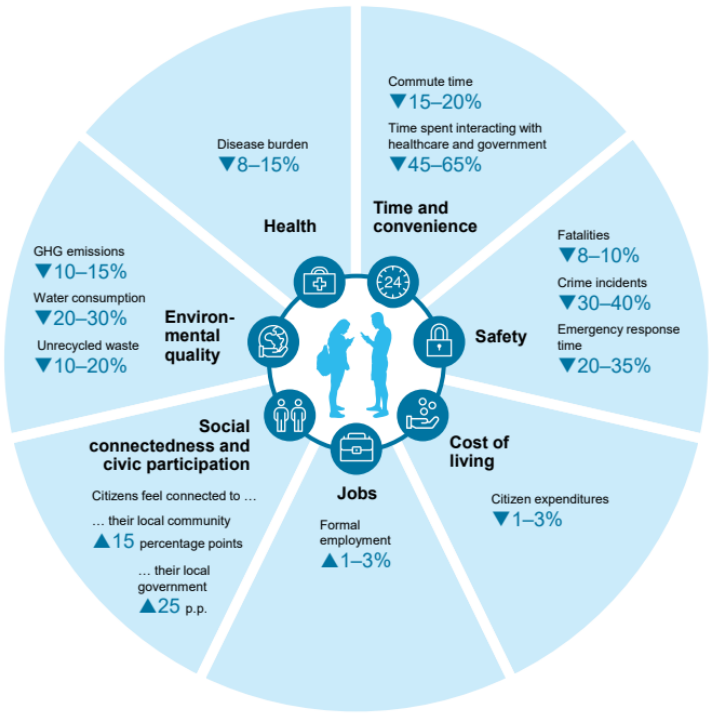

Figure 1 above illustrates the potential impacts and benefits of smart city AI applications across various dimensions, including:

1. **Safety**: Smart city applications can reduce fatalities by 8-10% and lower crime rates through real-time crime mapping, predictive policing, and the use of body-worn cameras by police.
2. **Time and Convenience**: Improvements in urban mobility through e-hailing services, real-time public transit information, and intelligent traffic systems can reduce commute times by 15-20%.
3. **Health**: The implementation of healthcare applications, including remote patient monitoring and digital scheduling, could reduce the disease burden by 8-15%.
4. **Environmental Quality**: Smart applications for energy management, such as smart grids and automated building systems, can cut greenhouse gas emissions by 10-15%.
5. **Social Connectedness and Civic Participation**: Digital platforms that enhance civic engagement, such as apps for community discussions and voting on city decisions, are being piloted in several cities to boost social participation.

Source: World Bank Group

VSaaS AI and Lenovo ThinkEdge Smart Cities Solutions

As cities grow and evolve, ensuring public safety becomes an increasingly intricate challenge. The Smart Cities solutions address this by harnessing the power of video analytics and IoT technologies. Lenovo and VSaaS.ai demonstrate how integrating various types of cameras, edge devices, and AI capabilities can create a robust public safety system.

VSaaS AI Solutions

As urban environments expand and become more complex, the need for innovative solutions to address various city challenges, from public safety to transportation efficiency, becomes paramount. Lenovo and VSaaS AI are at the forefront of this transformation with Smart Cities Solutions, integrating advanced video analytics, IoT technologies, and edge AI computing. These solutions not only enhance public safety but also optimize urban operations, improve business intelligence, and elevate the overall quality of life for residents. By leveraging a highly compatible and scalable platform, Lenovo and VSaaS AI offer cities a robust framework for future growth, setting a new standard for smart city initiatives. By incorporating edge devices and AI, we can:

- Enhance video quality and minimize latency
- Detect and promptly alert authorities to potential security threats
- Analyze traffic patterns to optimize transportation systems
- Provide business owners with actionable insights into customer behavior
- Offer residents a secure and convenient way to monitor their homes

Key Benefits of the Smart City Project:

- Improved Public Safety and Security
- Enhanced Transportation Efficiency

- Increased Business Intelligence and Revenue Growth
- Better Quality of Life for Residents
- Connect exiting camera CCTV infrastructure

This solution offers a scalable and adaptable platform, designed to support future expansion and innovation. By harnessing advanced video analytics and IoT technologies, the Smart City Project sets a new standard for public safety and smart city initiatives.

City managers often grapple with the challenge of integrating diverse CCTV infrastructures, including utility pole cameras, public transit cameras, small business cameras, and residential cameras. Edge computing serves as a bridge, connecting these complex and dispersed systems, enabling AI to detect anomalous activities and promptly send alerts to emergency control centers.

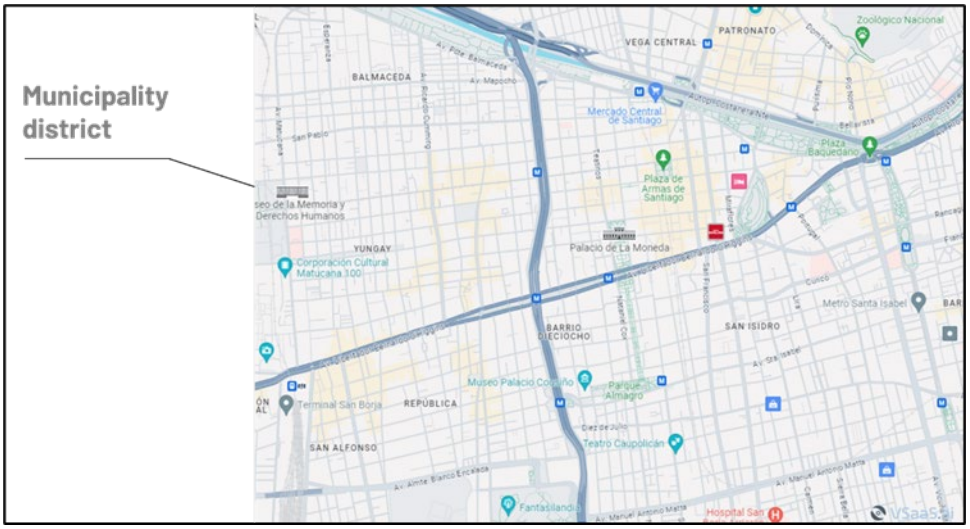

Figure 2 – City Map

Figure 3 – Public CCTV Cameras

Red markers are added to the map, representing the public CCTV cameras strategically installed across the city. These cameras serve as the backbone for monitoring public spaces and identifying potential security threats in real time.

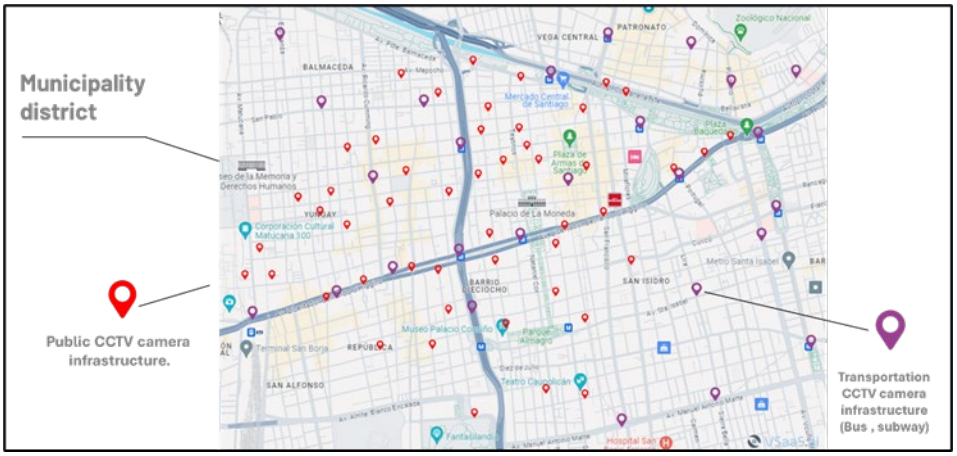

Figure 4 – Public Transportation Cameras

A layer of red markers is added to the map, symbolizing the public CCTV cameras strategically positioned throughout the city. These cameras establish a critical foundation for monitoring public areas and detecting potential security threats.

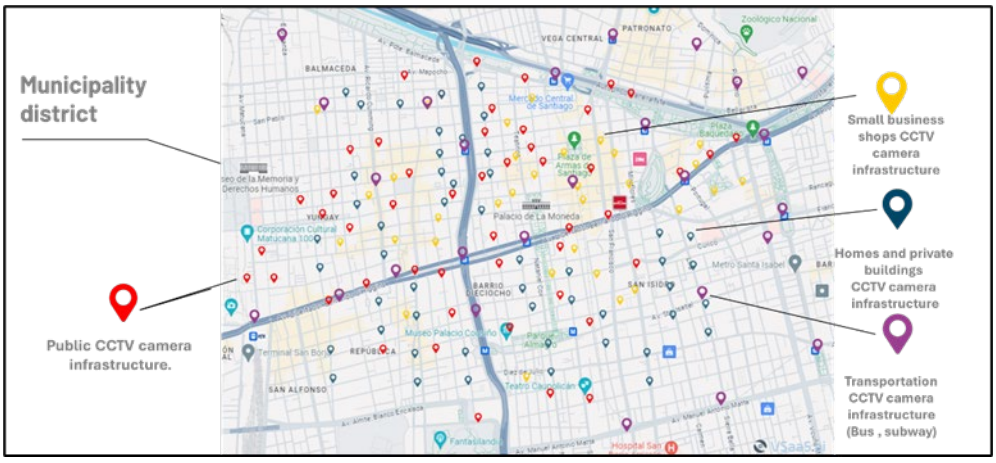

Figure 5 – Small Shops Cameras

A layer of yellow markers is introduced, representing cameras installed at small businesses and shops. These cameras play a crucial role in crime prevention and offer business owners valuable insights into customer behavior.

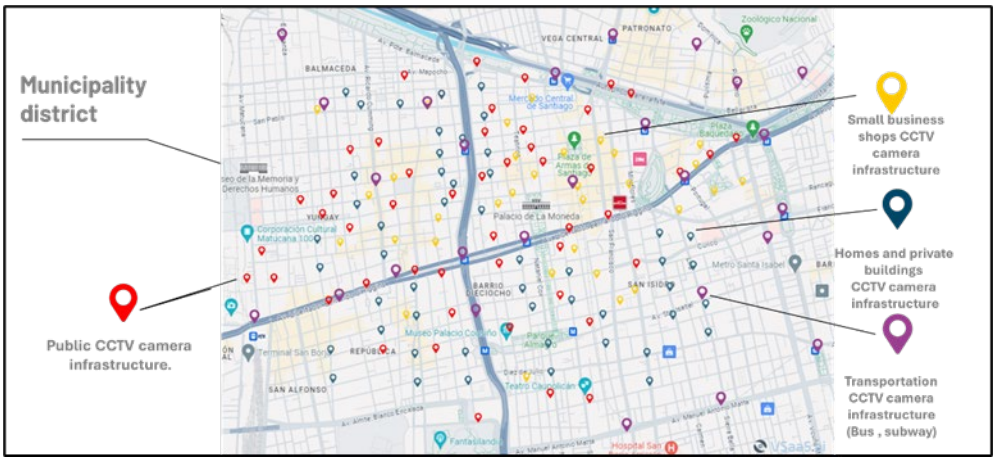

Figure 6 – Homes and Private Building Cameras

Finally, a layer of green markers is applied, indicating cameras installed at homes and private buildings. These cameras offer residents an enhanced level of security and peace of mind.

Edge AI Computing: Unlocking the Potential of Existing CCTV Infrastructure

Edge AI computing is a transformative technology that allows data processing and analysis to occur at the network's edge—closer to the data source. This approach minimizes latency, reduces bandwidth usage, and lowers costs associated with traditional cloud-based solutions. In the realm of CCTV infrastructure, edge AI computing can fully harness the capabilities of existing cameras, turning them into intelligent sensors capable of detecting, analyzing, and responding to events in real-time.

Benefits of Edge AI Computing for CCTV Infrastructure

- **Real-time Processing**: Enables immediate processing of video feeds, reducing latency and allowing for quicker response times.
- **Improved Accuracy**: Enhances the precision of event detection and decreases the occurrence of false positives by analyzing feeds in real-time.
- **Reduced Bandwidth**: Lowers the need to transmit video feeds to the cloud or central servers, thereby reducing bandwidth consumption and associated costs.
- **Enhanced Security**: Strengthens CCTV infrastructure security by enabling real-time threat detection and response.
- **Scalability**: Easily scales to support expanding CCTV infrastructure, making it ideal for large-scale deployments.
- Connecting Existing CCTV Infrastructure to Edge AI Computing

3. Technology Overview

This section provides a comprehensive exploration of the technologies deployed and the requirements addressed in VSaaS AI Smart Cities solutions.

The VSaaS AI platform offers a suite of AI-powered features designed to enhance the efficiency and effectiveness of video analytics systems.

- **Streamlined Camera Management:** Easily manage multiple cameras within a user-friendly, configurable environment.
- **Broad Compatibility:** Compatible with 99% of existing infrastructure, ensuring seamless integration.
- **Cost-Efficient Edge AI Computing:** Leverages edge AI computing to minimize reliance on centralized data centers, reducing operational costs.
- **Universal Access:** Enables connection and video streaming on any device, providing flexibility and convenience.
- **Interoperability:** Capable of working alongside other VMS systems, ensuring seamless parallel operation.

3.1 Functional requirements

A common use case for VSaaS involves clients who already have a CCTV infrastructure in place. For example, a municipality manages a network of approximately 2,500 cameras monitoring various neighborhoods across the city.

These camera feeds are sent to monitoring centers where operators observe events in real-time. However, this approach has limitations: it doesn't scale effectively, as operators cannot monitor every event simultaneously.

VSaaS addresses this challenge by leveraging the existing camera network to apply advanced analytics, filtering out irrelevant footage and highlighting only critical events in real-time. This allows operators to concentrate on fewer, more significant incidents, enabling them to respond more efficiently.

Figure ***1***. Data flows from the cameras to both the standard visualization system and the Operators Platform

The primary objectives of the solution are to:

- Deter criminal activity and detect suspicious behavior
- Reduce response times to incidents
- Enhance the overall sense of safety within neighborhoods
- Improve the availability and quality of evidence
- Optimize traffic management and coordination
- Implement and manage a blacklist of high-risk individuals or vehicles (License Plate Recognition blacklist)

Figure **2.** Monitoring Station: An operator utilizing VSaaS alert systems to manage approximately 350 cameras.

The main advantage of this distributed model is its extremely high fault tolerance. It also uses low bandwidth compared to streaming all the cameras' data into a central location, since only filtered events reach the platform for operators to see.

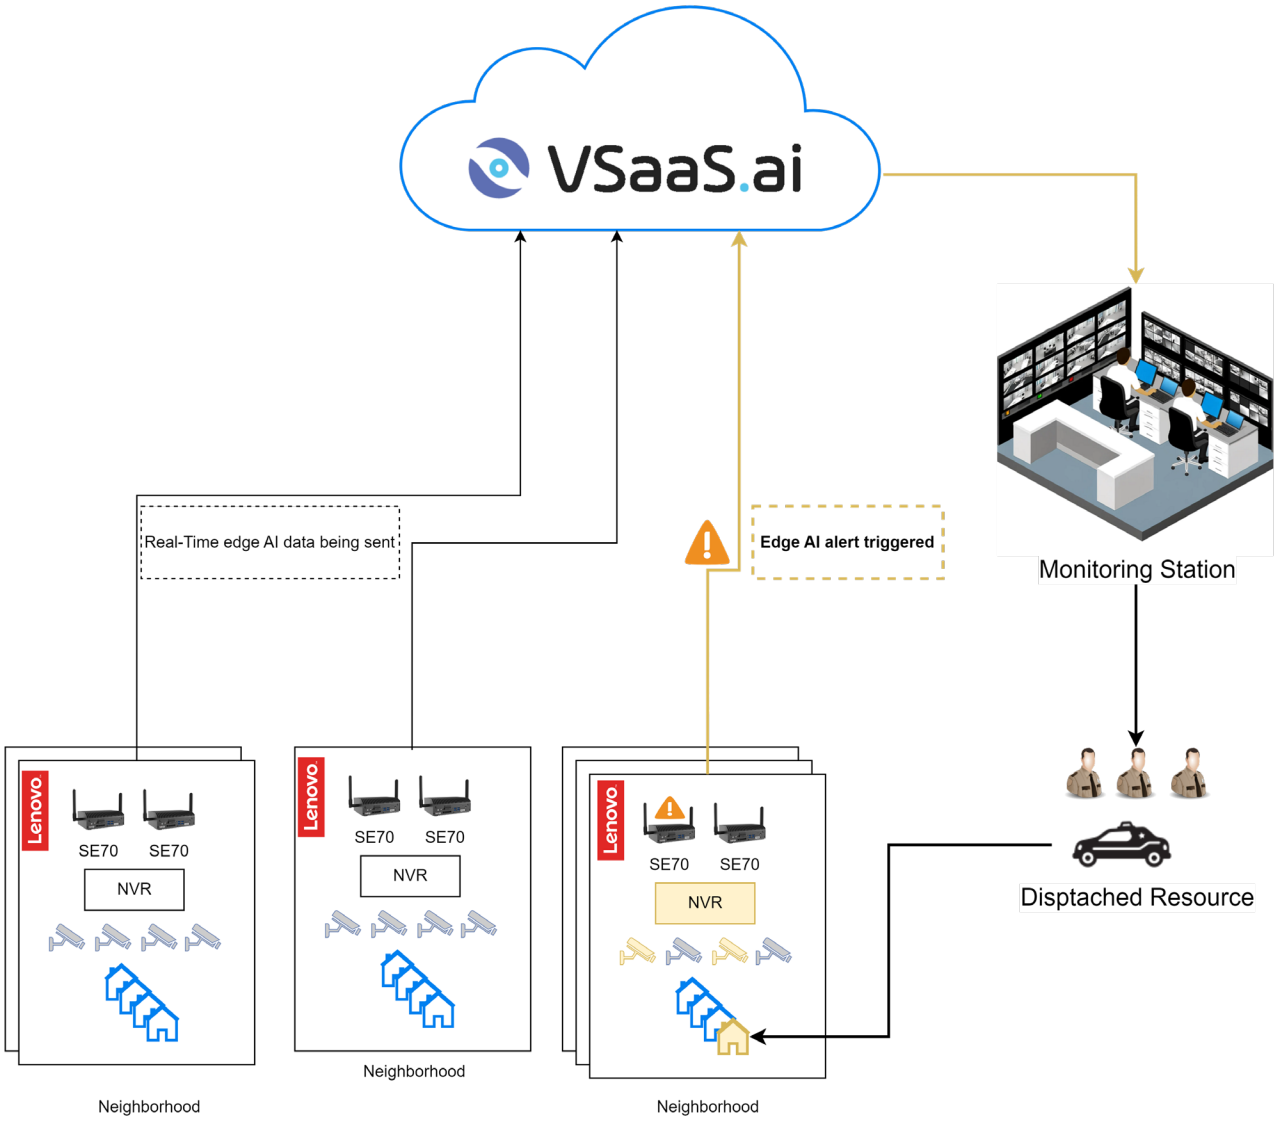

Figure 9. Monitoring Station layout concept.

3.1.1 Hardware Requirements

ThinkEdge SE70 Hardware Requirements

Infrastructure Requirements for Validation	1080p Cameras
ThinkEdge SE70 Deployment	6-core ARM v8.2
	8 GB DDR4 RAM
	512 GB NVMe Storage
	NVIDIA Jetson Xavier NX 16G.
	2x1 GB Ethernet
Interoperability	This validation focuses on hosting VSaaS AI applications on the platform. Clients can virtualize the environment to consolidate additional applications, but an interoperability test is required to ensure all ISV applications operate seamlessly on the same platform.
Performance Expectations/Threshold	Expected utilization of up to 80% of shared resources (i.e - CPU, RAM, HD). Maximizing Accelerator utilization (90-100%)

ThinkEdge SE360 V2 Hardware Requirements

Infrastructure Requirements	1080p Cameras
ThinkEdge SE360 V2 Deployment	16 Core Intel Xeon-D Processor (Hyper-thread Capable)
	256 GB of DDR4 3200 MHz RAM
	1 TB of NVMe Storage
	NVIDIA A2 Accelerator
	4x1 GB Ethernet
Interoperability	This validation focuses on hosting VSaaS AI applications. Clients can virtualize to add more applications. Interoperability tests are necessary to ensure seamless operation across all ISV applications on the platform.
Performance Expectations/Threshold	Expected utilization of up to 80% of shared resources (i.e -

	CPU, RAM, HD). Maximizing Accelerator utilization (90-100%)

ThinkEdge SE455 V3 Hardware Requirements

Infrastructure Requirements for Validation	1080p IP cameras
ThinkEdge SE455 V3 Deployment	32/64 Core AMD EPYC 8004 (Hyper-thread Capable)
	384 GB of DDR4 3200 MHz RAM
	1 TB of NVMe Storage
	NVIDIA L4/L40
	4x1 GB Ethernet
Interoperability	This validation focuses on hosting VSaaS AI applications. Clients can virtualize to add more applications. Interoperability tests are necessary to ensure seamless operation across all ISV applications on the platform.
Performance Expectations/Threshold	Expected utilization of up to 80% of shared resources (i.e - CPU, RAM, HD). Maximizing Accelerator utilization (90-100%)

3.1.2 IP Cameras

The VSaaS Smart City solution enhances safety, security, and operational efficiency by leveraging existing IP camera infrastructure for a wide range of applications. By utilizing current video monitoring or operational cameras, this solution enables edge AI-driven analytics directly at the data source, eliminating the need to transmit large volumes of data to central servers or the cloud. The solution's cameras, tested at a resolution of 1080p, ensure high-quality video capture, while the application processes video at a rate of 1-15 frames per second (fps). This configuration strikes a balance between real-time processing and data efficiency, making it particularly effective in scenarios where low latency and immediate response are essential.

3.1.3 - NVIDIA Accelerators

The NVIDIA Jetson platform is designed for edge AI, offering powerful AI capabilities in a compact form factor ideal for robotics and embedded systems. For energy-efficient edge deployments, the A2 GPU provides essential AI inference capabilities, making it a cost-effective solution for smaller-scale operations. The L4 accelerator strikes a balance between performance and efficiency, optimized for mainstream AI and video workloads in edge computing environments. At the top end, the L40 GPU delivers exceptional performance for demanding AI, graphics, and compute workloads, making it well-suited for advanced edge applications

requiring high scalability and real-time processing.

3.2 Non-functional requirements

Performance Requirements:

- **Real-time Processing**: The software must ensure rapid response times to user actions and system requests. The goal is to minimize latency, providing near-instantaneous feedback to enhance user satisfaction and efficiency. Typical response time targets should be defined based on specific use cases and industry standards.
- **Network Throughput Requirements:** The software should be capable of handling high volumes of data transfer across the network without bottlenecks. It must efficiently manage data transmission rates to support seamless communication and data exchange, ensuring that the network infrastructure can sustain peak loads and maintain performance levels. An average HD camera utilizes 1-5Mbps of bandwidth, depending on frame rate and bitrate. Also, the system must be capable of handling disconnection of edge devices and recover automatically.

Scalability Requirements:

- **Load Scalability:** The software must efficiently handle increasing loads without compromising performance. The system should scale horizontally to accommodate more users and data, ensuring consistent response times and throughput as demand grows.
- **Reliability**: The software is designed to operate without failures for a specified period under normal conditions. It should include mechanisms for error detection and recovery to minimize downtime and data loss, ensuring robust and uninterrupted service.
- **Availability Requirements:** The software must maintain a high level of availability, minimizing downtime and ensuring that the system is operational and accessible to users whenever needed. This includes implementing redundant systems and failover procedures to handle unexpected outages.
- **User Interface:** The software must offer a user-friendly interface that is intuitive and easy to navigate. The design should prioritize usability, incorporating clear visuals, consistent layouts, and rresponsive elements to enhance the overall user experience and ensure efficient interaction with the system.

Security Requirements

Data Encryption: The software must implement robust data encryption methods to protect sensitive information during transmission and storage. This includes using advanced encryption standards (e.g., SSL, AES-256, etc.) to safeguard data integrity and confidentiality, ensuring that user data is secure from unauthorized access and breaches.

Authentication: The software should provide secure authentication mechanisms to verify the identity of users accessing the system. This includes multi-factor authentication (MFA) and secure password policies to prevent unauthorized access, ensuring that only authorized users can access sensitive functions and data.

4. Design Overview

Lenovo Validated Design Architecture

This design illustrates the IT infrastructure supporting the VSaaS AI Smart Cities solutions. The system is engineered to manage a wide array of use cases; watchlist alerts, threat detection, intrusion detection, crowd forming, traffic analysis, etc.

The VSaaS application operates on top of a containerized microservices structure, which enables it to have a flexible architecture to better split load between the different devices, optimizing the available computing resources (CPU, Memory, Network, etc.). It relies mainly on video streams and snapshots coming directly from the Cameras and NVRs, but also can pull data from VMS systems.

This enables the architecture to adapt to the client's network, and work using hybrid systems easily to minimize the impact on network and physical resources, enabling easier scaling of the services.

The deployment options encompass both bare metal and cloud environments, with edge systems like SE360v2, SE450, and SE455v3 playing a crucial role. These edge servers are not only equipped with powerful compute capabilities to handle intensive processing tasks but are also designed with system resiliency in mind. They are built to withstand harsh environmental conditions, ensuring reliable performance even in challenging settings. This resilience is essential for maintaining consistent operations in diverse environments, making the VSaaS application robust and dependable.

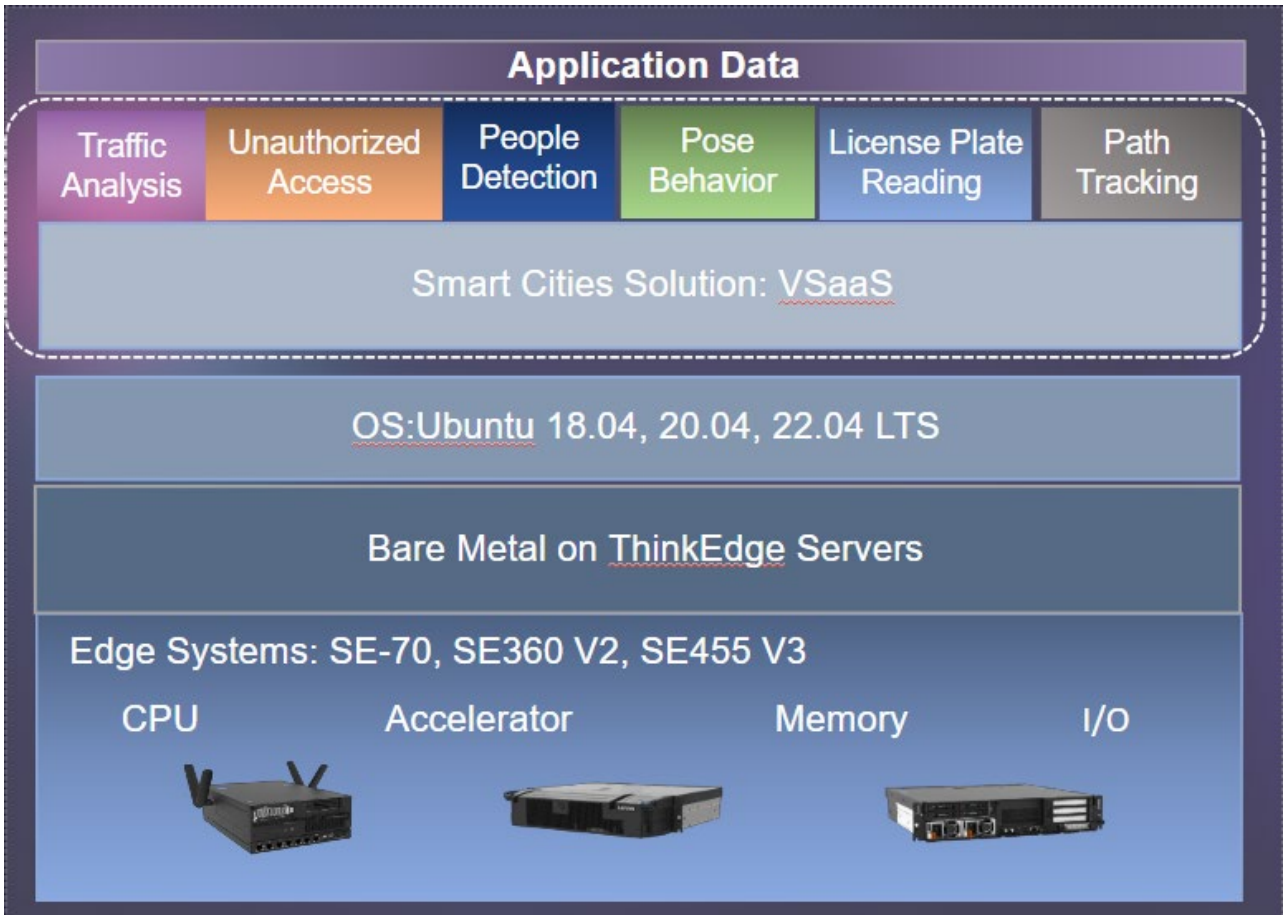

Figure 10. High level overview of the physical and logical components of this validated design

Key Value Proposition of VSaaS

One of the primary value propositions of VSaaS is its ability to adapt to existing architectures. With its proprietary orchestrator technology, AI workloads can be performed anywhere, enabling seamless integration with Lenovo's IoT family (SE10/30/50). The orchestrator efficiently sends video streams or images (fps) to the VSaaS cloud AI server for processing.

Example Use Case: SE10

See the image below, an example of how a SE10 can enable video analytics applications using VSaaS.

Edge Processing Capabilities

The VSaaS platform is also compatible with SE70, SE360, and SE455v3, which can connect multiple cameras and perform AI processing at the edge. The processed images are then sent to the VSaaS platform in the cloud for further analysis.

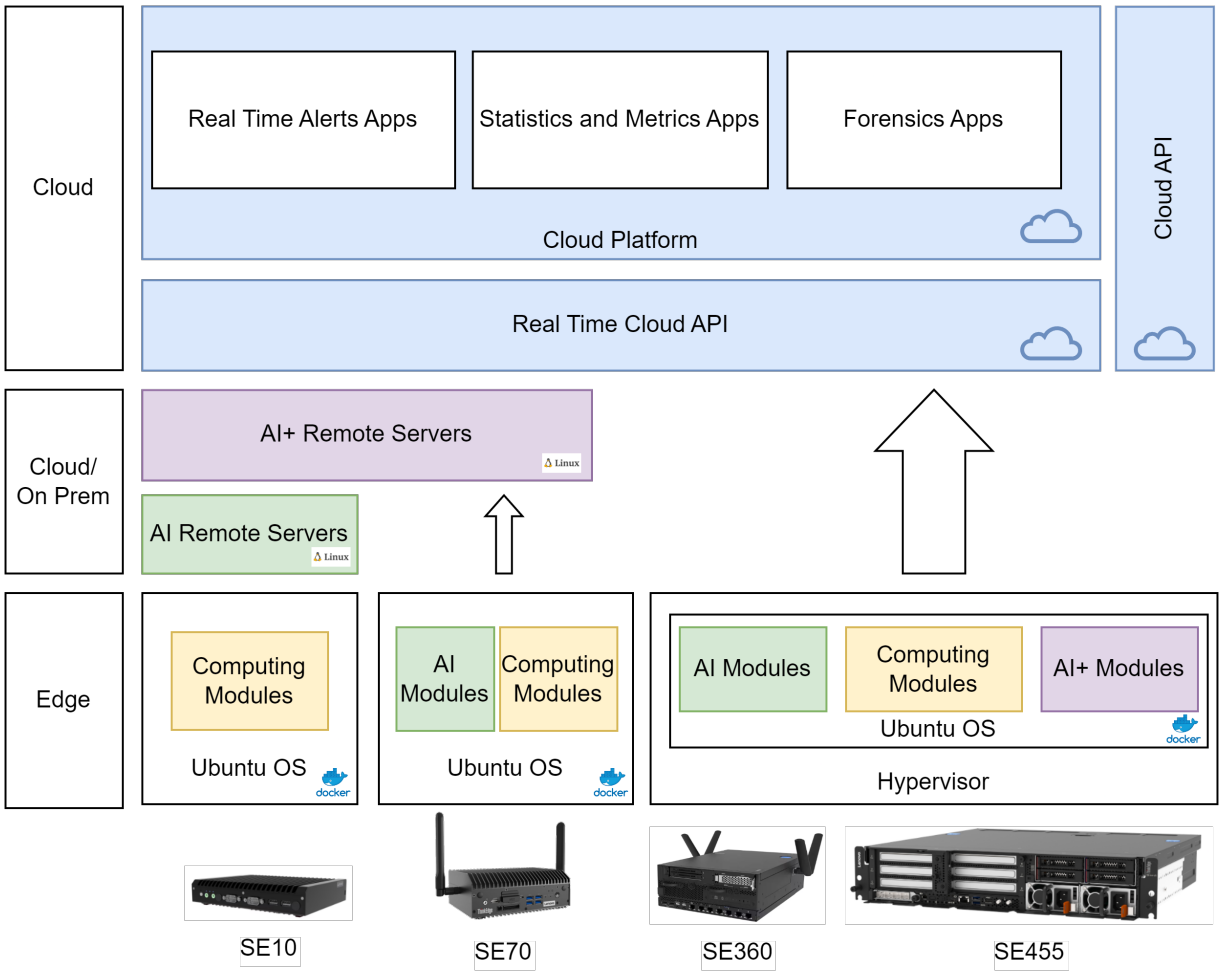

Figure 3. Software layers and hardware compatibility: Depending on device capabilities, various services can be processed at the Edge, on-premises, or in the cloud.

4.1 Component model

Connecting Existing CCTV Infrastructure to Edge AI Computing

- **Assess Existing Infrastructure:** Conduct a thorough evaluation of the current CCTV infrastructure, including the types of cameras, network architecture, and power supply to identify any compatibility or performance gaps.
- **Select Edge Devices:** Choose edge devices that are fully compatible with the existing infrastructure and are capable of processing video feeds in real-time to meet the specific needs of the system.
- **Install Edge Devices:** Strategically install edge devices either near the CCTV cameras or in a central location, based on the chosen system architecture, to optimize performance and coverage.
- **Configure Edge Devices:** Set up edge devices to efficiently process video feeds and detect events using advanced AI and machine learning algorithms tailored to the system's requirements.
- **Integrate with Existing Systems:** Seamlessly integrate edge devices with existing systems, such as video management software (VMS) and alarm systems, to create a unified and comprehensive security solution.

VSaaS AI Platform: A Catalyst for Edge AI Computing

VSaaS AI platform is instrumental in harnessing the full potential of edge AI computing within CCTV infrastructure. This platform can:

- **Manage Edge Devices:** Efficiently oversee and monitor edge devices, ensuring they operate optimally and process video feeds in real-time.
- **Provide AI and ML Capabilities:** Equip edge devices with advanced AI and machine learning capabilities, enabling them to detect and analyze events with high accuracy.
- **Seamlessly Integrate with Existing Systems:** Connect seamlessly with existing systems, such as video management software (VMS) and alarm systems, to deliver a comprehensive and cohesive security solution.
- **Ensure Scalability and Flexibility:** Provide a scalable and flexible framework, allowing users to easily add or remove edge devices as requirements evolve.

By integrating edge AI computing with a VSaaS AI platform, existing CCTV infrastructure is transformed into a powerful, intelligent, and responsive security system, setting a new standard for monitoring and public safety.

4.2 VSaaS Dashboard for Smart Cities

Traditional CCTV control rooms face operational challenges such as:

- Monitoring numerous cameras in real-time
- Handling different event types
- Integrating cameras from various brands
- Connecting multiple sites with diverse network architectures

VSaaS has developed a platform that enables a CCTV operator to manage over 300 cameras in real-time. Key features include:

- Managing customer accounts and apps
- Event management, forensic analysis, and license plate analytics tools
- Dashboards for CCTV events, count analytics, and Power BI or other BI integration
- Connecting and configuring edge devices and IP/analog cameras
- Setting up AI analytics and monitoring devices (edge, NVR, cameras)
- Remotely upgrading AI detection models
- API for third-party integration
- White labeling option available if required

Figures 12 – 20 showcase some of the different screens, modules and configurations available in the VSaaS application. For a more comprehensive run through of the platform please reach out to your regional sales representative.

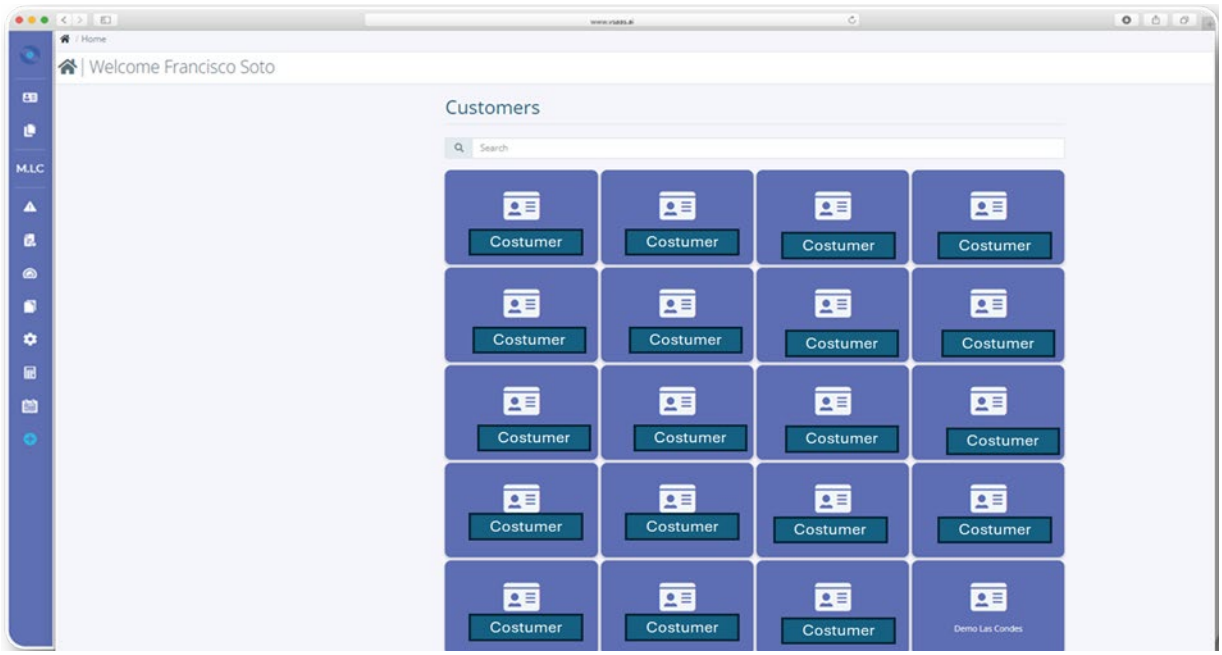

Figure 12. Management of costumer dashboard.

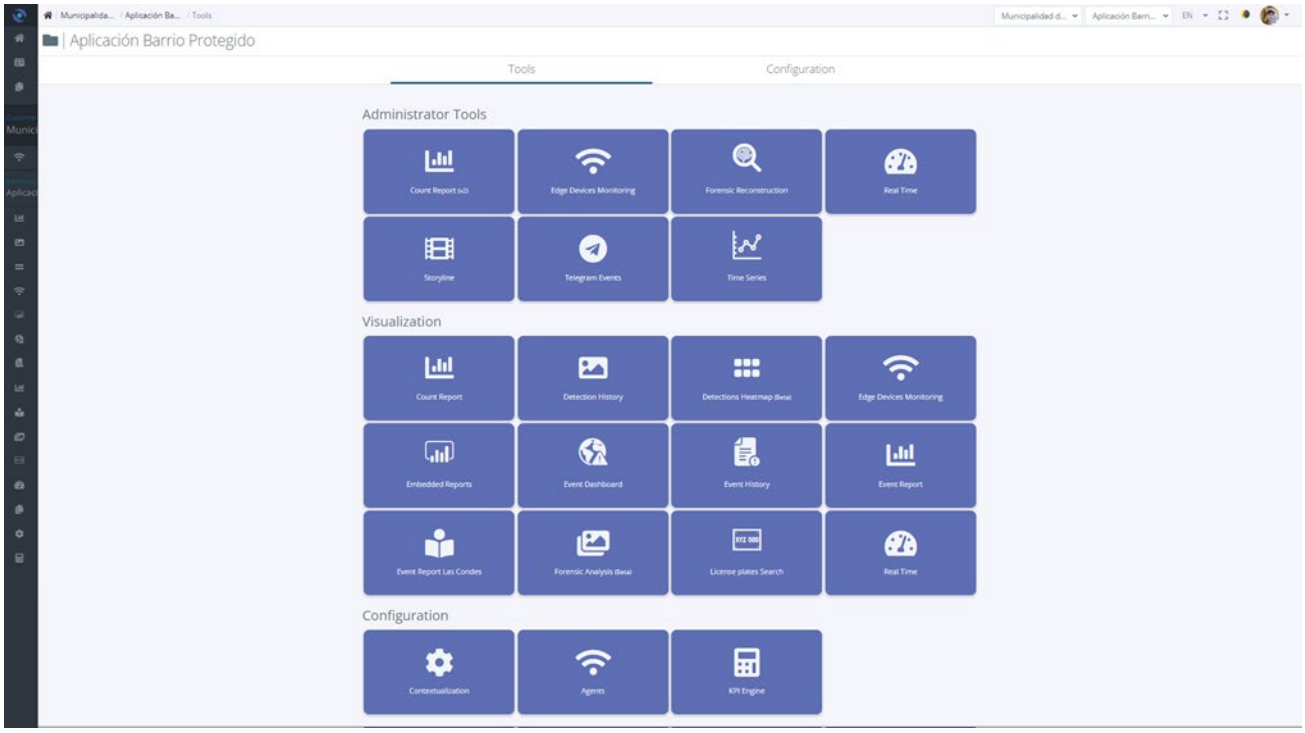

Figure 13. APP Tools. City cameras monitoring

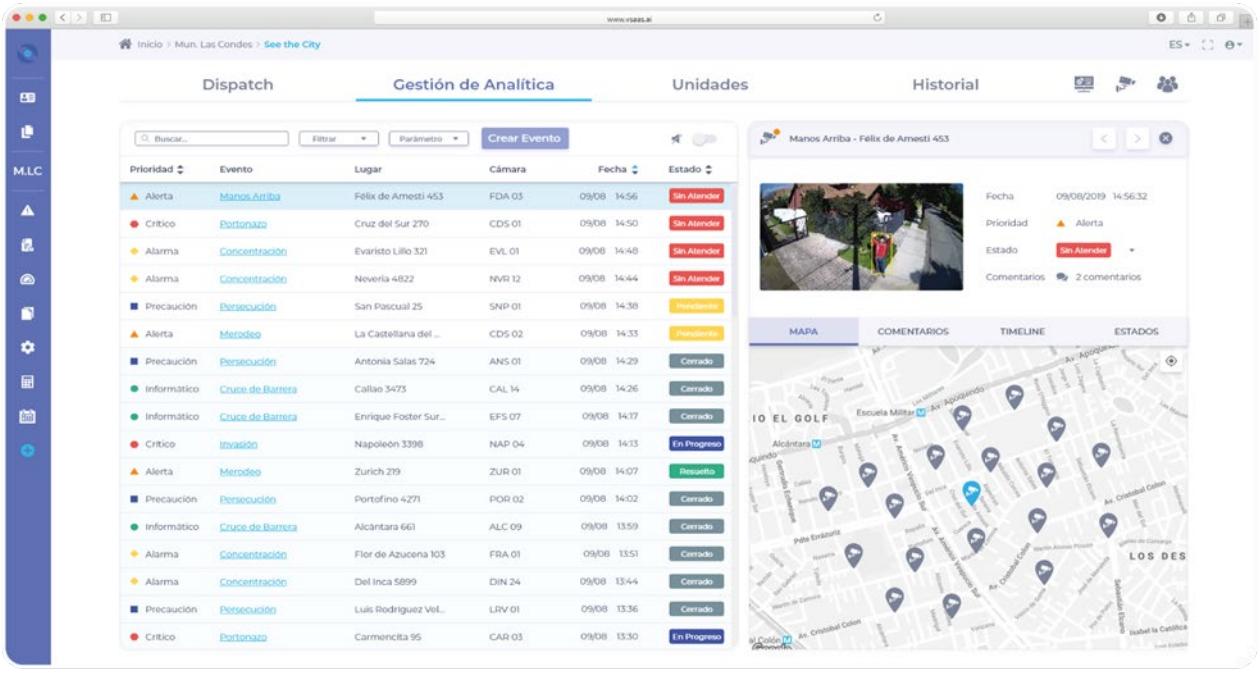

Figure 14. Tool event manager.

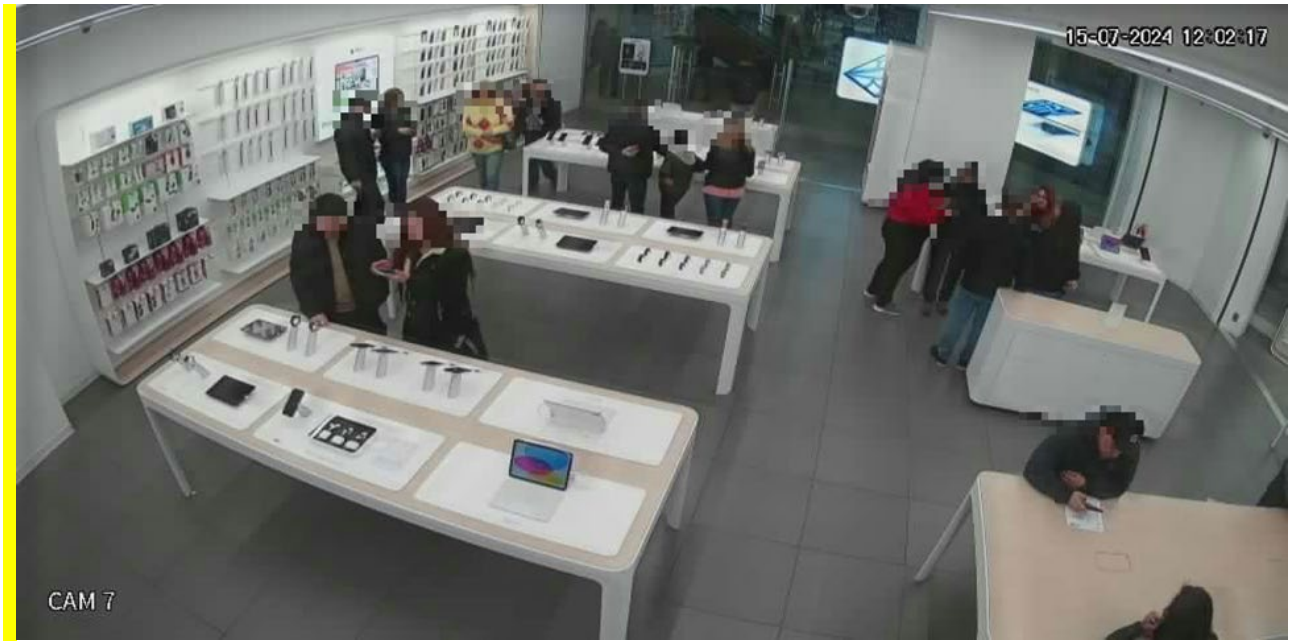

Figure 15. Blur faces of costumer in a store.

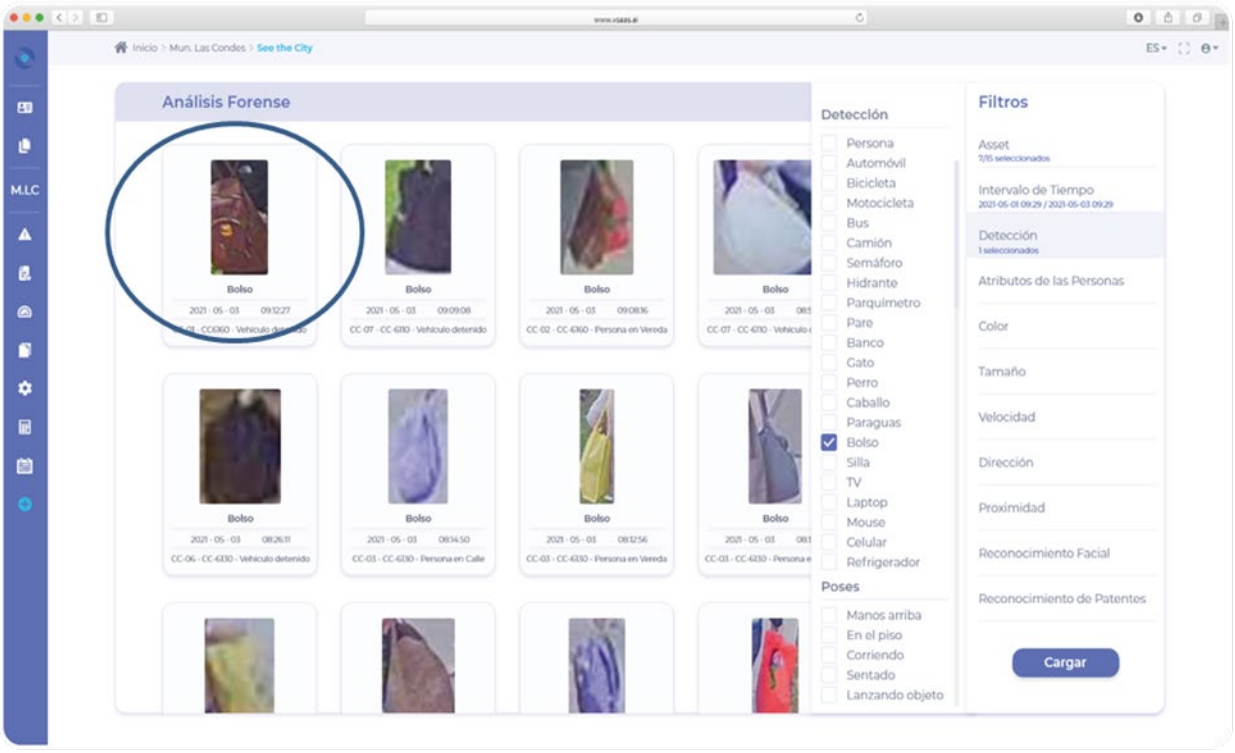

Figure 16. Forensic analytics Tool 1-2.

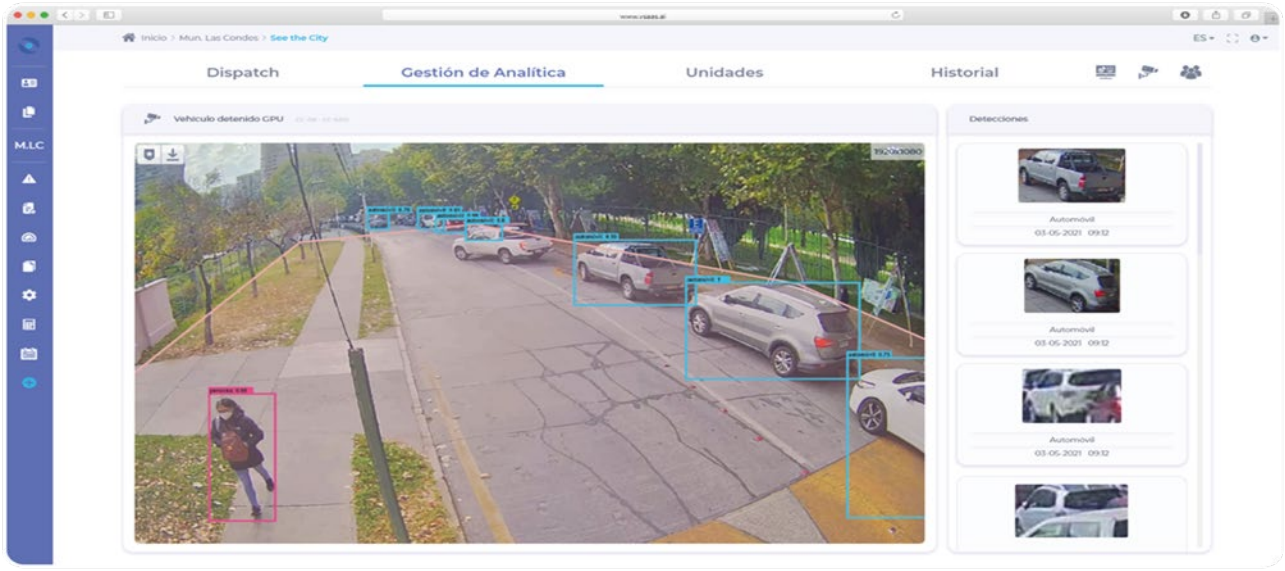

Figure 17. Forensic analytics Tool 2-2.

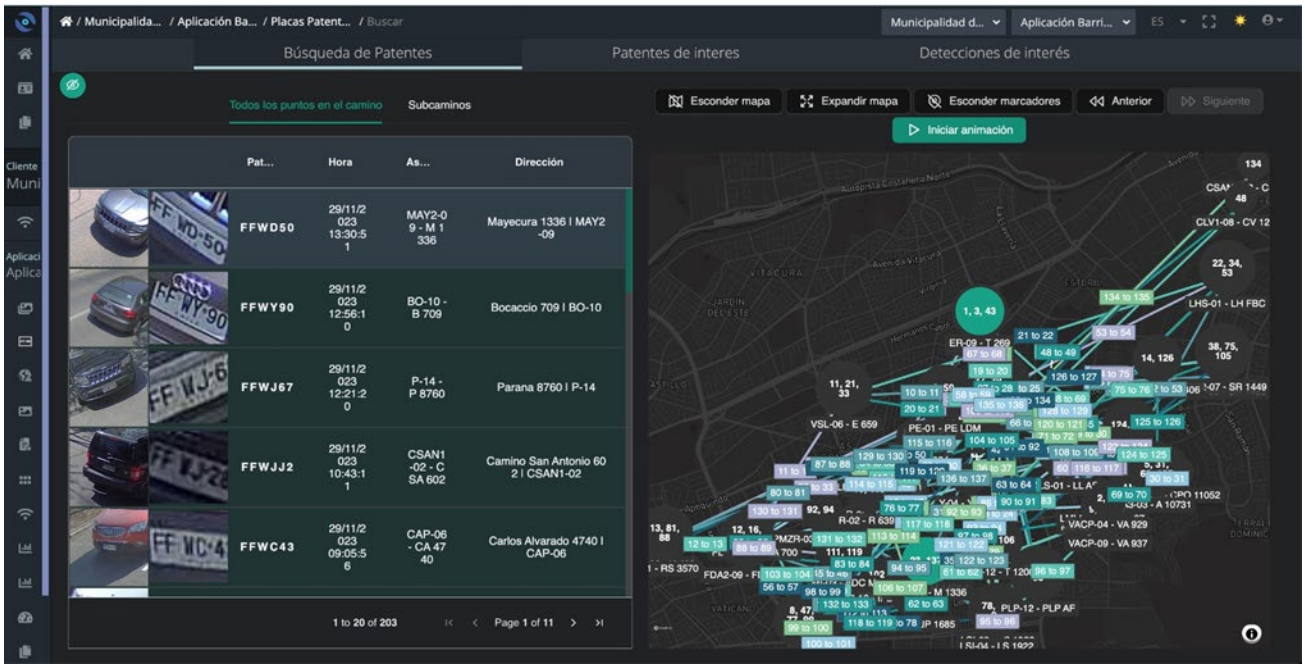

Figure 18. Tool forensic

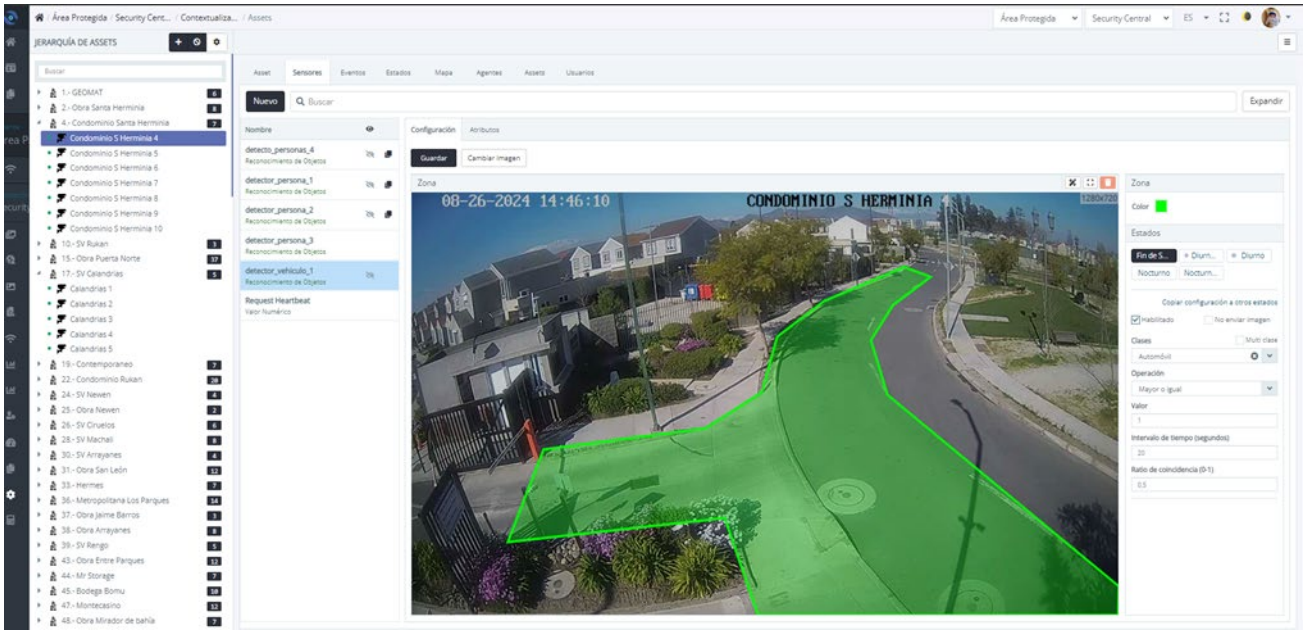

Figure 19. Analytics and edge configuration tool

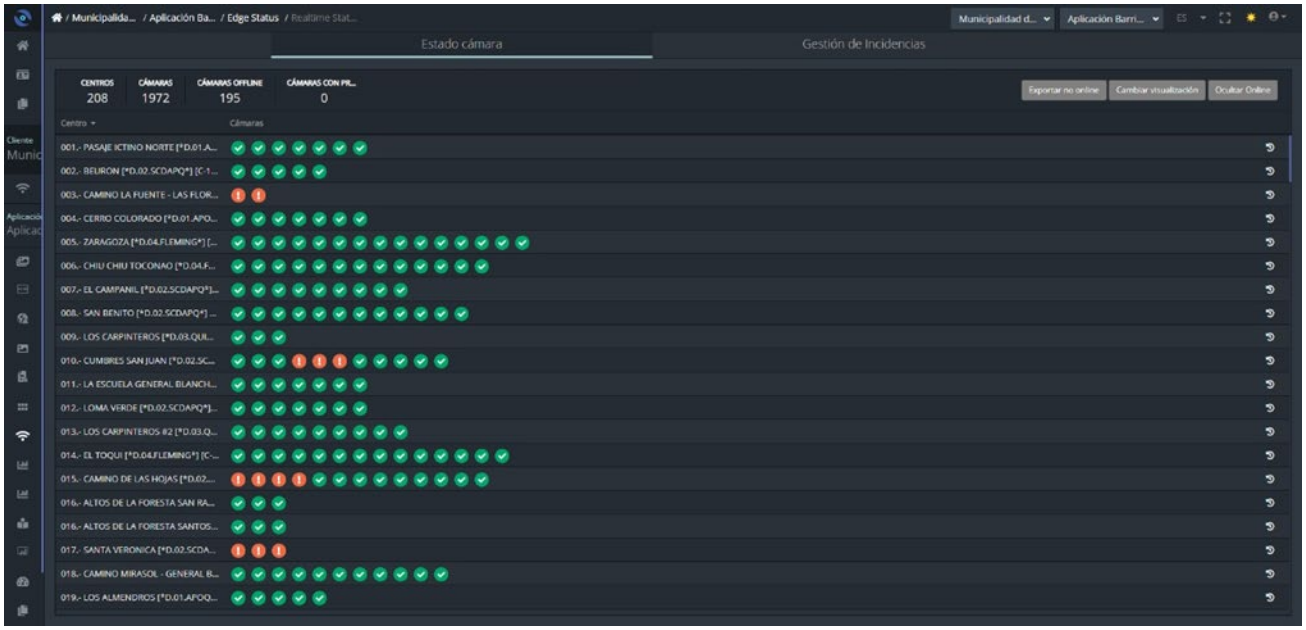

Figure 20. Cameras and edge monitoring Tool.

5. Deployment model

VSaaS AI compute workload can be deployed in multiple architecture modes: full edge-AI, hybrid (Edge-AI + Cloud-AI) and “gateway” (edge connector + Cloud-AI).

5.1. Full Edge Solution

In a fully edge-based environment, VSaaS video analytics software suite is deployed on SE70s within the same internal network as the NVR and cameras. All analytics are processed at the edge, with resulting events transmitted to the cloud.

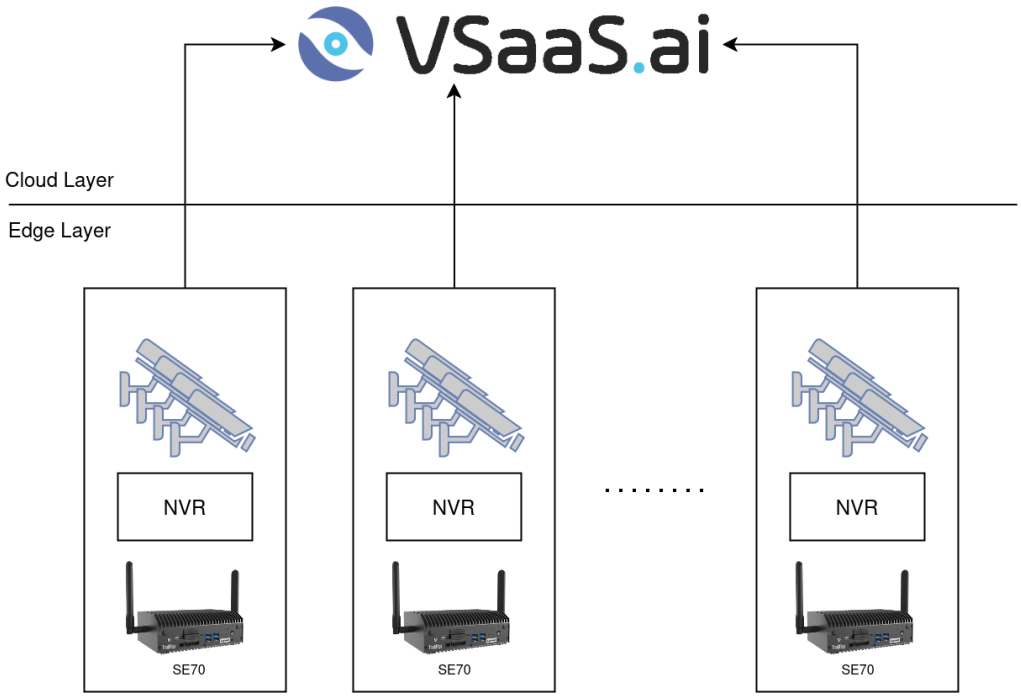

Figure 21. Full Edge Solution leveraging the SE70

5.2. Hybrid Solution

In a hybrid environment, VSaaS video analytics software suite is partially deployed on SE70/SE360/SE455V3 devices within the same internal network as the NVR and cameras, while advanced analytics are handled by a centralized compute cluster. Some analytics are processed at the edge, while those exceeding edge capacity are offloaded to the VSaaS Cloud-AI server platform.

This solution is ideal for environments with:

- Multiple locations
- 1-256 cameras per location
- Locations with 16-246 cameras can utilize various Lenovo ThinkEdge devices
- Cloud-based analytics, such as license plate recognition, human-pose detection, or color detection, deployed on VSaaS Cloud-AI servers

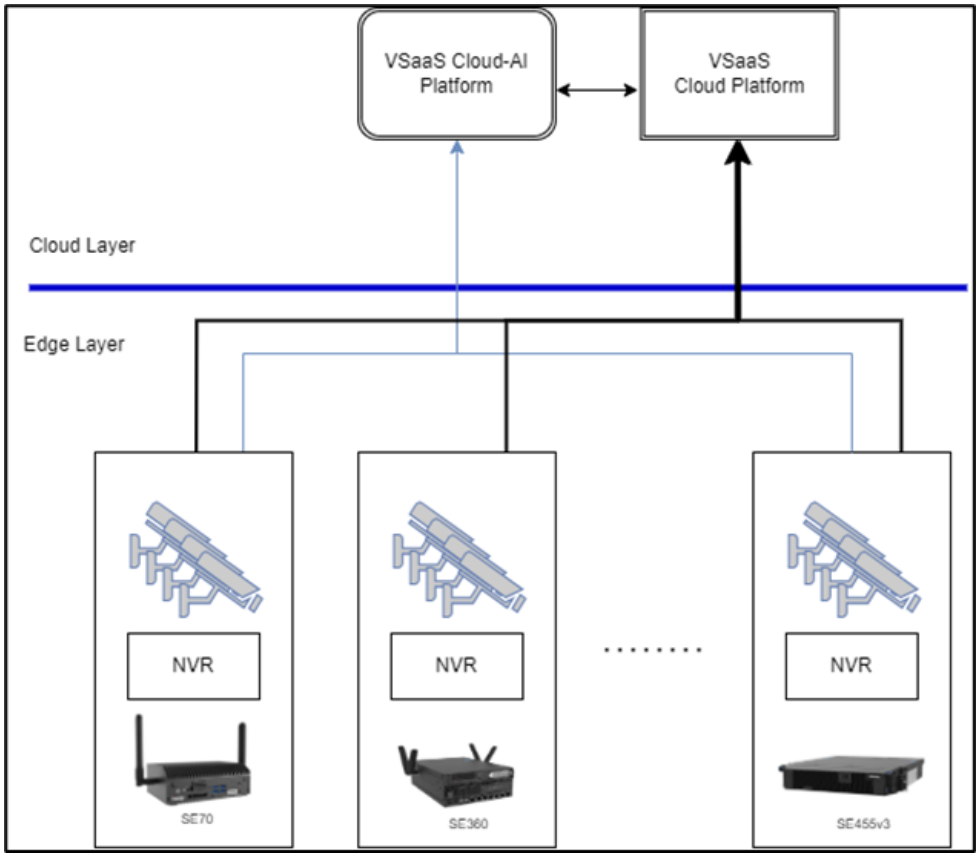

Figure 22. Hybrid Edge Solution leveraging the SE70, SE360 V2 & SE455 V3

5.3. Smart-Gateway Solution (Full AI in the cloud)

Customers increasingly prefer AI as a service, making it practical to deploy Lenovo edge devices (without AI acceleration cards) as smart gateways. These devices connect cameras to the VSaaS cloud AI server, extracting images and sending them to the cloud for analysis. Analytics are computed in the cloud, with events either pushed from the gateways or pulled by the cloud-AI servers.

This solution is ideal for environments where cameras/NVRs are behind firewalls, allowing centralized access via secure SSH tunnels or VPNs. It supports:

- Multiple locations with 1-32 cameras each
- Internet connectivity capable of 1 MB/s per camera
- Basic and advanced analytics deployed on VSaaS Cloud-AI servers

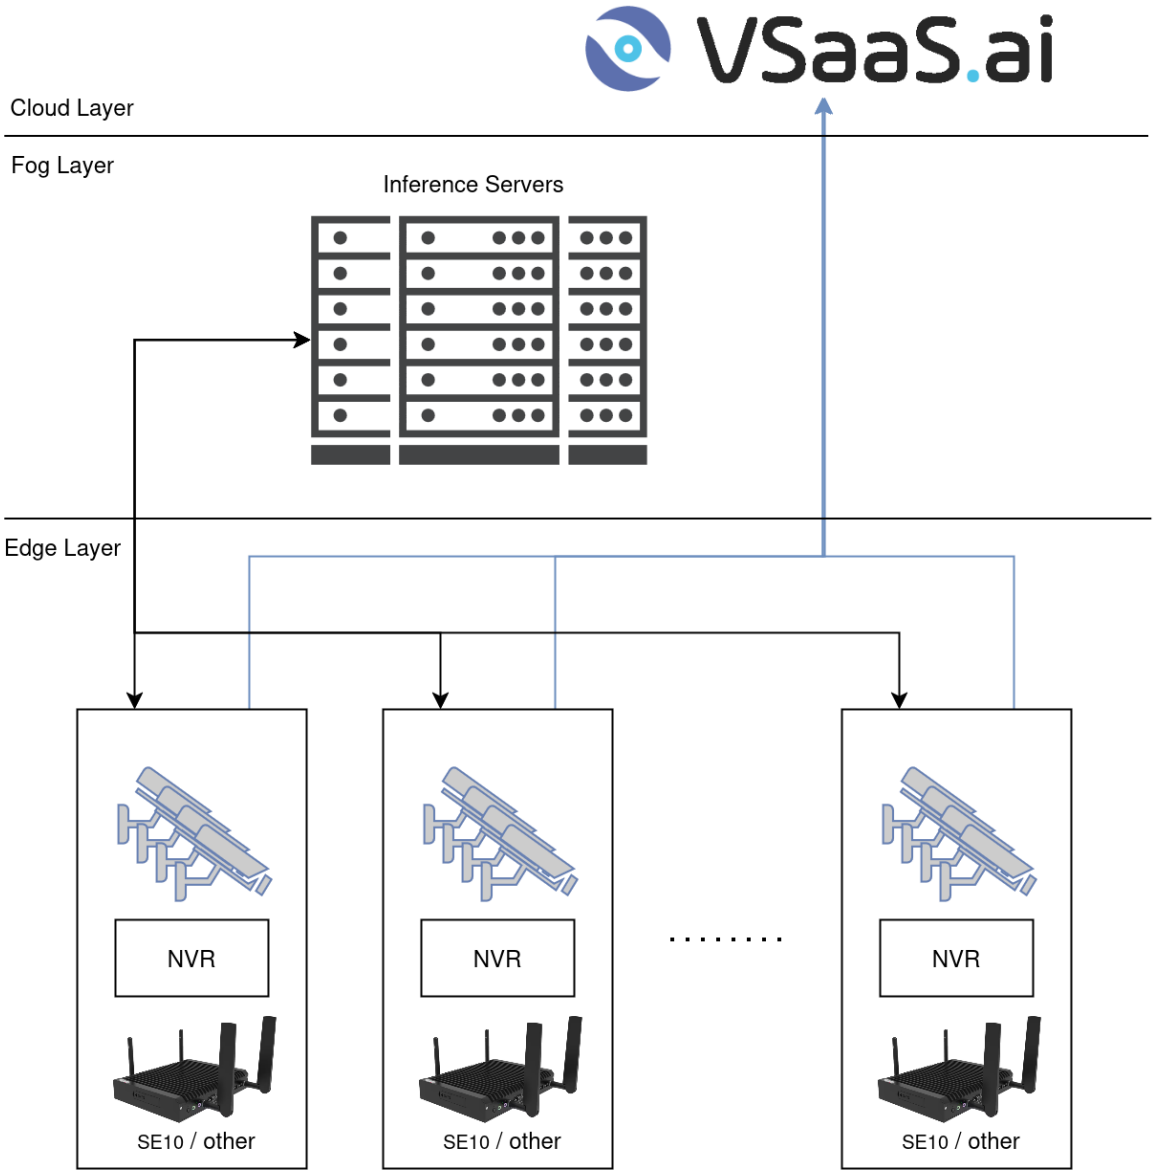

Figure 23. Smart-Gateway Solution Leveraging the SE10

6. Deployment considerations

Deploying VSaaS AI requires careful planning across several key areas to ensure optimal performance and integration within your business environment. These considerations include systems management, compute nodes, networking, and storage integration.

6.1 Systems Management

Effective systems management is essential for maintaining the VSaaS AI application. Implementing robust monitoring and logging to track system performance and quickly identify issues. Utilizing automated tools for updates, patches, and configurations to keep the application secure and up to date. Implementing access control and role-based management to protect sensitive data and system integrity. VSaaS runs on standard Ubuntu 20.04 LTS and can be managed using standard tools.

6.2 Server/Compute Nodes

Deploying VSaaS AI requires a well-planned server and compute node infrastructure. For high performance, use on-premises servers or scalable cloud-based solutions tailored for AI workloads. Ensure your compute nodes have sufficient CPU, GPU, and memory resources to handle intensive data processing and machine learning tasks. Incorporate redundancy and failover capabilities to maintain high availability and minimize downtime. VSaaS recommends dedicated bare metal servers running Ubuntu + Docker + Kubernetes for highest performance.

6.3 Networking

Networking is critical for data flow and application responsiveness in VSaaS AI deployment. Networking is a valuable resource of the organization that must be considered when implementing deployments for AI. Intensive networking consumption is not an option for many organizations, therefore leveraging the flexibility that the multiple Edge Computing options provide, can reduce the expected bandwidth consumption up to x20 times, enabling a whole new world of capabilities on top of the existing infrastructure.

6.4 Storage Integration

Efficient storage integration is vital for managing the large volumes of data processed by VSaaS AI. Choose high-performance SSDs (NVMe recommended) for fast data access and scalable cloud storage for long-term data retention. Organize and index data efficiently to facilitate quick retrieval and analysis. Implement data backup and disaster recovery plans to preserve critical information and enable restoration in case of failure.

By addressing these deployment aspects with precise and targeted strategies, businesses can ensure a smooth and successful integration of the VSaaS AI application, maximizing its benefits and capabilities.

6.5 Sizing & Licensing Guide

6.5.1 Sizing Guide

VSaaS and Lenovo have collaborated to develop a comprehensive sizing guide for compute resources, ensuring optimal performance and efficiency in VSaaS AI deployments. Accurate resource sizing is crucial for managing the intensive workloads of AI applications, helping you achieve the best possible performance and outcomes.

Lenovo Devices			Edge AI Solution	Hybrid AI Solution	Edge-AI Solution	Hybrid AI Solution	Smart-Gateway
Model	Proposed HW	AI CARD					
Edge AI	SE10	n/a	n/a	Up to 32 cameras			x
Edge AI	SE30	n/a	n/a	Up to 64 cameras			x
Edge AI	SE50	n/a	n/a	Up to 128 cameras			x
Edge AI	SE70	Jetson NX	Up to 8 cameras	Up to 32 cameras	x	x	
Edge AI	SE360 V2	A2	Up to 16 cameras	Up to 64 cameras	x	x	
Edge AI	SE455 V3	2 x L4	Up to 96 cameras	Up to 512 cameras	x	x	
Edge AI	SE455 V3	2 X L40	Up to 340 cameras	Up to 512 cameras	x	x	
Edge AI	SE455 V3	n/a	n/a	Up to 1024 cameras			x

Table 4: Sizing Guide

6.5.2 VSaaS Features and Use Cases Guide

Enhanced Analytics Capabilities with Lenovo Devices

A Lenovo ThinkEdge Server can deploy analytics on both edge and cloud AI servers, leveraging the AI capabilities of the VSaaS infrastructure. This allows more complex analytics, beyond the edge's capacity, to be seamlessly distributed to AI cloud servers using VSaaS's proprietary orchestration technology.

Lenovo Model	Edge AI Solution		Hybrid AI Solution		
	Standard	AI+ Tracking	Standard	AI-Cloud LPR	AI-pose
SE10/30/70			x	x	x
SE70	x	x		x	x
SE360 V2	x	x		x	x
SE455 V3	x	x		x	x

Table 5: Lenovo ThinkEdge models and VSaaS license types (Edge and Cloud)

1. Standard Features:

- Object detection and storage for forensic analysis
- Detection of objects in specific zones and monitoring their presence over time
- Alerts based on combined conditions of object presence and duration in zones
- Configurable settings for different times and dates
- Daily operational reports for alerts
- Advanced rule creation for objects, pedestrians, workers, time, and zones

Sample Object Types: Pedestrian, Car, Bicycle, Motorcycle, Bus, Truck, Cat, Dog, fire Hydrant, trash bag, Cone, Barrier, garbage collector, Scooter, Backpack, food delivery Backpack, shopping cart, cell phone, helmet, shoe box, chair, door, forklift, excavator, bulldozer, flatbed-trailer. Streetlight, car with No-plate, face detections.

Use Case Examples:

- Detecting people in prohibited areas
- Monitoring after-hours activities for people or vehicles
- Identifying unauthorized movement in safety zones
- Detecting vehicles in restricted areas
- Alerting when a person or vehicle lingers too long
- Identifying unauthorized workers in restricted zones
- Monitoring objects like bags, boxes, or materials on the floor
- Ensuring Personal Protective Equipment (PPE) compliance
- Detecting caution signs (e.g., wet floor)
- Monitoring gate/door activity
- Tracking occupancy, dwell time, and movement
- Arrival, delivery, and dispatch tracking

- Weapon detection
- Detecting camera tampering

2. **AI+ Tracking Features:**

- Counting people and objects (cars, motorcycles, buses, trucks, scooters)
- Vehicle traffic analysis
- Traffic light color detection

Use Case Examples:

- Heatmapping and traffic analytics
- People counting and dwell time measurement
- Crowd flow analysis (e.g., wrong-way detection)
- Detecting vehicles running red lights
- Drive-thru vehicle identification, counting, and analysis

3. AI-Cloud License Plate Recognition (LPR)/ Optical Character Recognition (OCR)

Near Real-Time LPR and OCR:

- License plate reading and storage
- Blacklist alerting within 10-30 seconds
- OCR reading (e.g., container text)

License Plate Forensic and OCR (Text Detection):

- License plate reading and storage (detections saved over 60 minutes)
- Blacklist alerting after 60 minutes
- OCR reading (e.g., container text)

Use Case Examples:

- Track vehicle movement across multiple cameras for forensic analysis
- Retrieve event detection history for a vehicle
- Alert when a blacklisted vehicle enters the facility
- Drive-through vehicle identification
- Analog instruments reding

4. **AI-Pose Features:**

- Hands-up, fall, crouching, jumping detection

Use Case Examples:

- Detecting a person on the floor
- Identifying fare evasion through 'back-cocking'

[1] Q2 2024. VSaaS update new model quarterly.

6.6 Importance of Sizing Recommendations

Accurate sizing is essential to ensure that your infrastructure can handle the intensive data processing and AI inferencing tasks required by VSaaS AI. Following the guidelines developed by VSaaS and Lenovo helps businesses avoid pitfalls like under-provisioning, which can cause performance issues, and over-provisioning, which can lead to unnecessary costs. The sizing recommendations below serve as a general guideline, but each deployment may require adjustments to meet specific needs.

VSaaS AI applications require a robust and well-balanced compute environment. The sizing guide specifies requirements for CPUs, GPUs, and memory to efficiently manage diverse AI workloads. High-performance CPUs are essential for general data processing, while powerful GPUs are crucial for training and inferencing machine learning models. Sufficient memory is also vital to ensure smooth operation and rapid data access.

7. Validation Scope

Lenovo validation design framework helps ensure that the infrastructure is optimized for AI-driven VSaaS smart cities applications, featuring industry-leading servers, software, networking, and storage. The validation scope primarily covers:

- **Speed and Reliability:**
  The designs enable rapid deployment of edge-optimized applications, minimizing the time needed for research, design, and troubleshooting.
- **Customized Solutions:**
  Flexible architecture allows for tailored adaptations to meet specific deployment needs, enhancing efficiency and effectiveness.
- **Comprehensive Testing:**
  Lenovo and third-party partners rigorously test and validate the design, reducing proof-of-concept time and eliminating lengthy configuration processes.

Validation Scope:

The validation scope involves Lenovo hardware and VSaaS application focuses on ensuring that this solution meets specific performance and interoperability standards suitable for various deployment scenarios.

Performance Testing:
Assesses how Lenovo hardware handles VSaaS applications under typical and extreme conditions, including:

- **Load Testing:** Ensures the system performs well under expected and peak loads without degradation.
- **Stress Testing:** Pushes the system beyond normal capacities to evaluate its ability to handle extreme workloads and manage errors.

Interoperability Testing:
Ensures Lenovo hardware integrates seamlessly with VSaaS software, focusing on:

- **Compatibility Testing:** Verifies that the software operates smoothly on Lenovo hardware across various environments without conflicts.
- **Integration Testing:** Confirms seamless data flow across the platform, with correct API, data format, and protocol alignment.
- **Network Interoperability:** Tests the performance and reliability of network connections, crucial for cloud-based or distributed systems.

This validation process provides businesses with confidence that their investment in Lenovo hardware and VSaaS applications will result in a robust, efficient, and compatible IT environment, ready to meet operational demands.

8. Test and Validation Methodology and Results

The VSaaS Smart Cities solution was validated in a Lenovo lab environment to ensure accuracy and reliability. Key variables such as camera resolution, network speed, and deployment model were carefully managed. The validation used 1080p cameras, ingesting video, depending on the model, at a rate of 1-15 fps. Optimized network conditions ensured efficient AI-driven analytics. Use cases included person identification, threat detection, and object detection.

The number of concurrent cameras streams the system can handle may vary based on different use cases. Factors such as camera types, network bandwidth, and deployment models at the customer's site can influence performance, affecting the system's effectiveness and responsiveness in real-world deployments.

The testing methodology for the **VSaaS** application was designed to rigorously evaluate its capabilities in object detection, threat detection, and facial recognition, people, and automotive tracking. We meticulously validated the application on various servers, including the **SE70**, **SE360v2**, and **SE455v3**, to ensure comprehensive coverage and accurate assessment across different hardware configurations.

To simulate real-world scenarios, we used test RTSP streams that mimicked the behavior of IP camera streams. This approach provided a controlled yet realistic environment, ensuring our tests closely reflected real-world conditions.

This validation will focus on the following key areas:

Dynamic Use Cases:

- **People Detection:** Evaluate the system's accuracy and responsiveness in identifying and tracking individuals across various environments.
- **License Plate Tracking:** Assess the precision and reliability of license plate recognition and tracking under different conditions and camera angles.
- **Body Pose Estimation:** Analyze the system's capability to detect and interpret human body poses, with potential applications in security and behavior analysis.
- **Additional Use Cases:** Explore other functionalities supported by VSaaS.ai, such as facial recognition, object detection, and threat detection, traffic analysis, etc. ensuring they meet performance benchmarks in diverse scenarios.

Infrastructure and Hardware Integration:

Lenovo ThinkEdge Servers:

- **SE70:** Validate the performance of entry-level edge computing in low-power, space-constrained environments.
- **SE360 V2:** Assess the mid-range capabilities for handling AI workloads in environments with moderate power and space requirements.
- **SE455 V3:** Evaluate the high-performance edge computing capabilities in large-scale, demanding environments.

NVIDIA Accelerators:

- **Jetson:** Measure the effectiveness of edge AI acceleration for real-time processing tasks.
- **A2, L4, L40:** Validate the scalability and performance improvements provided by these accelerators in managing complex AI models and workloads.

By leveraging the SE70, SE360 V2, and SE455 V3 servers, we tested the VSaaS application's performance under various conditions and configurations. This thorough testing process validated the application's performance, reliability, and scalability.

Validated Results

Table 6 below highlights the validation components. Our collaborative testing demonstrated that the application delivered outstanding performance on the SE70, SE360v2, and SE455v3 servers, excelling in object detection, threat detection, license plate recognition (LPR), and person and vehicle tracking. Moreover, the efficient utilization of compute resources met our high expectations and standards. For more detailed insights, please contact your sales representative.

Table 6 – Validated configuration

HW Configuration		Cameras: AI on the Edge		HW Utilization				
Proposed HW	AI CARD	AI Standard (1 fps @ 1080p)	AI+ Tracking (10 fps @ 1080p)	CPU %	RAM %	GPU %	Network	Storage %
SE70	Jetson NX	8	2	Please contact your regional Sales Representative for detailed performance data.				
SE360 V2	1 x A2	16	1					
SE455 V3 (64 Core)	2 x L4	96	6					
SE455 V3 (32 Core)	2 X L40	340	22					

Appendix: Lenovo Bill of materials

ThinkEdge SE360 V2 with NVIDIA A2

7DAMCTOAWW	ThinkEdge SE360 V2 with Controlled GPU	1
BS56	ThinkEdge SE360 V2 Chassis	1
BS58	ThinkEdge SE360 V2 4x 1Gb, 2x 2.5Gb(TSN) I/O Module	1
BW2H	System Operational Temperature 5C to 40C / 41F to 104F	1
BFYE	Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
BS42	ThinkEdge SE350 V2/SE360 V2 Planar with Intel Xeon D-2775TE 16C 100W 2.0 GHz	1
B966	ThinkSystem 64GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	4
5977	Select Storage devices - no configured RAID required	1
AVV0	On Board SATA Software RAID Mode	1
BC4V	Non RAID NVMe	1
BS5M	ThinkEdge SE360 V2 M.2 Cabled Adapter Module	1
BQUJ	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	2
BS2C	ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU	1
BS5J	ThinkEdge SE360 V2 Riser Assembly (PCIe Riser + 7mm Backplane)	1
BUGM	Internal AC Power Supply (230V/115V)	1
BUGU	ThinkEdge SE360 V2 AC Power Input Board	1
BW8U	ThinkEdge SE360 V2 500W 230V/115V Non-Hot Swap Power Supply	1
BS5W	ThinkEdge SE360 V2 Fan Assembly (Front to Rear)	1
BYBQ	XClarity Controller Managed	1
B0MK	Enable TPM 2.0	1
BRPJ	XCC Platinum	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

ThinkEdge SE455 V3 – 32 Core with NVIDIA L40

7DBYCTOAWW	ThinkEdge SE455 V3 - 3Yr Warranty with Controlled GPU	1
BVTK	ThinkEdge SE455 V3 Chassis	1
BFYA	Operating mode selection for: "Maximum Efficiency Mode"	1
BY8S	System Operational Temperature 5C to 25C / 41F to 77F	1
BW2T	ThinkEdge SE455 V3 AMD EPYC 8324P 32C 180W 2.65GHz Processor	1
BW3M	ThinkEdge SE455 V3 2U Heatsink	1
BUVV	ThinkSystem 96GB TruDDR5 4800MHz (2Rx4) 10x4 RDIMM-A	4
5977	Select Storage devices - no configured RAID required	1
BVUU	ThinkEdge SE455 V3 2.5" NVMe Backplane	1

BVUY	ThinkEdge SE455 V3 M.2 SATA/x4 NVMe Adapter with Carrier	1
BQUJ	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	2
BN2T	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-Port OCP Ethernet Adapter	1
BS2C	ThinkSystem NVIDIA L40 48GB PCIe Gen4 Passive GPU	2
BVUR	ThinkEdge SE455 V3 Riser1	1
BVUS	ThinkEdge SE455 V3 Riser2	1
BMH8	ThinkEdge 1100W 230V/115V Platinum Hot-Swap Power Supply	2
BYBQ	XClarity Controller Managed	1
BPKR	TPM 2.0	1
BVV6	ThinkEdge SE455 V3 Intrusion Switch	1
BVTX	ThinkEdge SE455 V3 Standard EIA Bracket	1
BVTL	ThinkEdge SE455 V3 Motherboard	1
BRPJ	XCC Platinum	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

ThinkEdge SE455 V3 – 64 Core with NVIDIA L4

7DBYCTOAWW	ThinkEdge SE455 V3 - 3Yr Warranty with Controlled GPU	1
BVTK	ThinkEdge SE455 V3 Chassis	1
BFYA	Operating mode selection for: "Maximum Efficiency Mode"	1
BY8S	System Operational Temperature 5C to 25C / 41F to 77F	1
BW2R	ThinkEdge SE455 V3 AMD EPYC 8534P 64C 200W 2.3GHz Processor	1
BW3M	ThinkEdge SE455 V3 2U Heatsink	1
BUVV	ThinkSystem 96GB TruDDR5 4800MHz (2Rx4) 10x4 RDIMM-A	4
5977	Select Storage devices - no configured RAID required	1
BVUU	ThinkEdge SE455 V3 2.5" NVMe Backplane	1
BVUY	ThinkEdge SE455 V3 M.2 SATA/x4 NVMe Adapter with Carrier	1
BQUJ	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	2
BN2T	ThinkSystem Broadcom 57414 10/25GbE SFP28 2-Port OCP Ethernet Adapter	1
BS2C	ThinkSystem NVIDIA L4 24GB PCIe Gen4 Passive GPU	2
BVUR	ThinkEdge SE455 V3 Riser1	1
BVUS	ThinkEdge SE455 V3 Riser2	1
BMH8	ThinkEdge 1100W 230V/115V Platinum Hot-Swap Power Supply	2
BYBQ	XClarity Controller Managed	1
BPKR	TPM 2.0	1
BVV6	ThinkEdge SE455 V3 Intrusion Switch	1
BVTX	ThinkEdge SE455 V3 Standard EIA Bracket	1
BVTL	ThinkEdge SE455 V3 Motherboard	1
BRPJ	XCC Platinum	1

7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

Appendix: Abbreviations

Acronym	Meaning
ISV	Independent Software Vendor
OS	Operating System
GUI	Graphic User Interface
AI	Artificial Intelligence
LVD	Lenovo Validated Design
CPU	Centra Processing Unit
GPU	Graphic Processing Unit
RAM	Random Access Memory
VM	Virtualized Machine
VMS	Video Management System
IP Camera	Internet Protocol Camera
RTSP	Real-Time Streaming Protocol
VPN	Virtual Private Network
GDP	Gross Domestic Product
SSH	Secure Shell Protocol
CCTV	Closed-Circuit TV
LPR	License Plate Recognition
OCR	Optical Character Recognition
NVR	Network Video Recorder

Resources

Resources	Links
SE360 V2	Lenovo ThinkEdge SE360 V2 Server Product Guide > Lenovo Press
SE455 V3	Lenovo ThinkEdge SE455 V3 Server Product Guide > Lenovo Press
VMware Horizon View	vmware.com/products/horizon-view
VMware vSphere	Vmware.com/products/data center -virtualization/vsphere
World Bank Group	Urban Development (worldbank.org)
VSaaS AI	https://www.vsaas.ai/
Smart Cities: Digital Solutions for a More Livable Future	https://www.mckinsey.com/capabilities/operations/our-insights/smart-cities-digital-solutions-for-a-more-livable-future#part2
VSaaS.ai case studies	Keeping the streets safer with AI at the edge | Lenovo US
Smart City of Barcelona	https://www.lenovo.com/us/en/case-studies-customer-success-stories/smart-city-of-barcelona0

Document history

Version 1.0	August, 2024	First version includes Lenovo ThinkEdge and VSaaS AI Smart Cities solutions.
Version 1.1	October, 2024	This version includes an updated BOM to reflect sizing recommendations.

Trademarks and special notices

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®
ThinkEdge®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

AMD and AMD EPYC™ are trademarks of Advanced Micro Devices, Inc.

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Power BI is a trademark of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo’s current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.