Lenovo Validated Design: Smart Cities AI for Public Safety with Vaidio

Last update: **26 September 2025**

Version 1.0

AI Video Intelligence with 30+ analytics for increasing safety and operational efficiency

Improve traffic efficiency with video analytics AI agents

Pre-Tested Design, deployment guides, sizing, and validation to speed up time to deployment

Optimized performance and scalability on Lenovo Servers

Vanita Meyer

Jorge Alberzoni

Lumir D'Ambrosio

LENOVO PRESS

Table of Contents

Introduction 1

City Challenges and Opportunity 2

City Challenges to Solve 2

Opportunity 2

Technical Overview 4

Functional Requirements 4

Video Ingestion & Analytics 4

Non-functional Requirements 5

Architectural overview 7

Operational Model 11

Key Operational Aspects 11

Vaidio Operational Model 11

NVIDIA Technology 12

Deployment considerations 13

Systems management 13

Server / Compute Nodes 14

Lenovo ThinkEdge SE100 Edge Server 14

Networking 16

Storage integration 16

Solution Validation Results 17

Validation Scope 17

Validation Results 17

Solution Summary 19

Appendix A: Vaidio Analytics Capabilities 20

Appendix B: Lenovo Bill of materials (BOM) 21

Appendix C: Abbreviations 26

Resources..... 27

Document history..... 28

Trademarks and special notices 29

Introduction

This Lenovo Validated Design (LVD) delivers public safety solutions in collaboration with the Vaidio, an AI Vision Platform, built on Lenovo ThinkEdge infrastructure. Vaidio transforms camera-based visual data into real-time intelligence with 30+ analytics applications, including anomaly detection, vehicle and license plate recognition, and incident alerts. Designed for cities, the Vaidio solution improves situational awareness, traffic management, and emergency response, while also supporting campuses, healthcare, retail, and enterprise use cases.

With automated setup, sub-2-second search, and deployment across edge, cloud, or on-premises platforms, Vaidio helps cities scale from pilot projects to city-wide rollouts. Validated on Lenovo SE100 and SE455 V3 servers with NVIDIA GPU acceleration, this LVD ensures performance, interoperability, and efficiency. Deployment guides, benchmarks, and sizing models reduce risk, speed time-to-value, and provide a repeatable blueprint for safer, smarter cities.

Intended Audience

This document is intended for enterprise IT decision-makers, solution architects, and technical professionals who sell into cities and involved in safety, security, or smart infrastructure initiatives. It is also relevant for Lenovo partners, systems integrators, and public sector leaders evaluating turnkey AI solutions to maximize usefulness of legacy video systems, enhance situational awareness, enable quick decision making, and improve response times across various environments such as cities, campuses, transportation hubs, and enterprise facilities.

City Challenges and Opportunity

City Challenges to Solve

Cities face mounting challenges in optimizing efficiency on roadways, safeguarding public spaces, and ensuring residents' quality of life. Traditional video monitoring systems rely on manual observation, which is costly, labor-intensive, and prone to error. With thousands of cameras deployed across intersections, transit hubs, and public areas, operators struggle to monitor streams in real time, detect incidents proactively, and respond quickly.

At the same time, rapid urban growth creates new pressures such as traffic congestion, pedestrian safety risks, and the need for compliance with environmental and security regulations. Yet many municipalities remain constrained by siloed technologies and legacy VMS/NVR systems that lack AI-driven intelligence, interoperability, and scalability.

Key limitations of current systems include:

- **Manual monitoring gaps:** Human operators cannot continuously watch large volumes of video feeds, resulting in missed events and delayed responses.
- **Scalability challenges:** Expanding coverage introduces operational overhead, integration complexity, and inconsistent performance.
- **Inefficient Search:** Reviewing hours of footage is slow, resource-intensive, and often inaccurate.
- **False alarm fatigue:** Legacy motion-detection systems generate excessive false positives, diverting attention from real incidents.

These limitations prevent cities from transforming video data into actionable intelligence, slowing response times and acting proactively, ultimately reducing situational awareness. The result is heightened risk to public safety, regulatory compliance, and continuity of urban operations.

Opportunity

The combined Vaidio and Lenovo ThinkEdge solution transforms video networks into intelligent, real-time platforms that improve safety, mobility, and operational efficiency. By automating monitoring and analysis, municipalities reduce reliance on manual supervision, lower costs, and scale safety operations with AI.

Proactive detection of accidents, traffic, pedestrian safety, and suspicious activities accelerates emergency response which improves mobility in congested areas and helps prevent incidents before they escalate. Compliance and safety analytics strengthen regulatory enforcement and reduce liability risks, while automated tools cut case resolution times from hours to minutes.

At scale, the solution enhances situational awareness, builds resident's trust, and supports safer streets, smoother transportation, and more resilient city services. Lenovo's validated design ensures predictable performance, allowing cities to start small and expand seamlessly with confidence.

This LVD solution delivers:

- **Scalability & Future-Readiness:** Edge-to-cloud architecture supports growth from localized pilots to city-wide rollouts, with validated small, medium, and large deployment models.
- **Improved Public Safety:** Automated detection of anomalies reduces response times and supports proactive prevention.
- **Smarter Traffic & Mobility Management:** Vehicle counting, license plate recognition, and congestion analytics optimize traffic flow.
- **Operational Efficiency:** Centralized monitoring and AI-driven alerts scales operator efficiency, enabling smaller teams to manage large networks.
- **Interoperability & Integration:** Standards-based design integrates with existing cameras command centers, and third-party platforms.
- **Hybrid, Open Architecture:** Kubernetes microservices support flexible deployment across edge, core, and cloud environments; ONVIF compliance and integration with 35+ VMS, PSIM, and sensor platforms ensure investment protection.
- **Accelerated AI Performance:** NVIDIA GPU acceleration delivers real-time analytics across thousands of streams, enabling sub-second search, faster incident detection and decision making.

By combining Vaidio's intelligent analytics with Lenovo's high-performance infrastructure, cities will experience enhanced public safety while creating sustainably efficient processes.

Technical Overview

The requirements for the Vaidio AI Vision Platform are divided into two categories**:** Functional Requirements and Non-Functional Requirements. Functional requirements define the specific capabilities and behaviours the system must deliver, such as video ingestion, object detection, search and alerts. They describe what the system should do to meet user and business needs across diverse use cases, including traffic management, public safety, and large-scale event monitoring. Non-functional requirements, on the other hand, describe the quality attributes and operational constraints of the solution. These include scalability, performance, reliability, security, usability, maintainability and integration with third-party systems ensuring that the system not only functions correctly but also operate efficiently, securely, and at scale.

Analyzing video is compute-intensive, particularly in complex configurations, where performance depends on the type and number of analytic engines enabled. Vaidio's functionalities are grouped into the following categories:

- **Vehicle** – Includes all transportation modes such as cars, trucks, motorcycles, bicycles, buses, and trains. Vehicle recognition supports traffic monitoring, autonomous driving, and urban planning.
- **People** – Anonymized detection is vital for security and situational awareness as sensitive road users.
- **Animal** – Covers both domestic and wild animals. Useful for wildlife monitoring, veterinary care, and pet use cases.
- **Object** – Refers to general inanimate items outside other categories, including furniture, electronics, tools, or bags. Plays a key role in inventory management and smart environments.

Functional Requirements

Video Ingestion & Analytics

- Support ingestion of Real Time Streaming Protocol (RTSP) streams and uploaded footage from any IP camera or Video management systems (VMS).
- Detect, classify, and track vehicles and objects in real time.
- Provide analytics functions, including pedestrian safety, license plate recognition, traffic analytics, accident detection, crowd detection and others.
- Enable cross-camera tracking across multiple feeds.
- Detect specialized objects (e.g. PPE compliance, smoke/fire).

For a complete list of supported analytics, see Appendix A – Vaidio Analytics Capabilities.

Search

- Allow search across live and recorded streams using object attributes and metadata.
- Deliver sub-2 second search for large-scale deployments.
- Support smart tags and hashtags for quick retrieval of specific events.
- Generate heatmaps and statistical reports over defined timeframes.

Alerts & Notifications

- Configure real-time alerts based on object detection, abnormal activity, or scene changes.
- Deliver alerts via multiple channels (email, HTTP/HTTPS, VMS/NVR integrations, mobile apps).
- Support cooldown intervals between repeated alerts to reduce alert fatigue.
- Provide acknowledgment workflows for alerts (Unclaimed, Claimed, Cleared).

Dashboards & Reporting

- Offer customizable dashboards for live monitoring and statistics.
- Support role-based access for dashboards and reports.
- Provide centralized management of thousands of cameras and nodes via Command Center.
- Include interactive visualization (charts, graphs, maps, floor plans, video walls).

Use Cases

Detailed use cases mapped to analytics are provided in Appendix A – Vaidio Analytics Capabilities.

Non-functional Requirements

Integration & Deployment

- Integrate with 30+ VMS and Network Video Recorders (NVRs)
- Deploy flexibly across on-premises, cloud, hybrid, and Kubernetes clusters.
- Expose APIs/SDKs for third-party integration (e.g. traffic management).
- Provide auto-configuration of cameras using AI for faster setup.

Performance & Scalability

- Scale to support 100,000+ cameras across distributed sites.
- Return search results within 2 seconds at scale.
- Support jumbo upgrades across versions without reinstallation.
- Maintain accuracy above 95% in real-world deployments.

Reliability & Availability

- Automatically restore camera states after upgrades or restarts.
- Provide redundancy in Core/Enterprise deployments via Kubernetes orchestration.
- Ensure uninterrupted analytics by decoupling Command Center management from Core nodes.

Security & Privacy

- Support API / HTTPS for VMS/NVR integrations.
- Provide role-based access control for dashboards, reports, and alerts.
- Blur/anonymize personally identifiable information by default; allow unblurred access only to authorized users.
- Comply with AI transparency and data traceability standards.

Usability

- Provide a web-based GUI accessible via any-browser.
- Offer intuitive dashboards with customizable layouts and visualizations.

- Reduce alert fatigue through configurable cooldown intervals.

Maintainability & Extensibility

- Allow addition of new AI models and analytics without major reinstallation.
- Support false detection reporting to continuously improve AI models.
- Enable integration of custom-trained analytics (e.g., domain-specific objects).

Efficiency & Optimization

- Optimize compute usage by dynamically loading lightweight vs. heavy AI models only when needed.
- Reduce deployment time by 90% through automated camera configuration.
- Minimize hardware load by offloading management functions to Command Center.

Architectural overview

The Vaidio AI Vision Platform architecture is designed to deliver scalable, secure, and high-performance video analytics across edge, core, and enterprise environments. It integrates seamlessly with existing camera networks, VMS/NVR platforms, and third-party applications, while leveraging Lenovo ThinkEdge infrastructure to provide real-time AI processing close to the data source. The architecture balances local edge intelligence with centralized orchestration, ensuring low latency, high availability, and the flexibility to support deployments ranging from individual sites to city-wide networks. This layered approach, spanning ingestion, analytics, management, and integration, provides the foundation for reliable and extensible AI-powered video monitoring.

Figure 1 shows the Vaidio AI Video Analytics solution architecture used for inferencing.

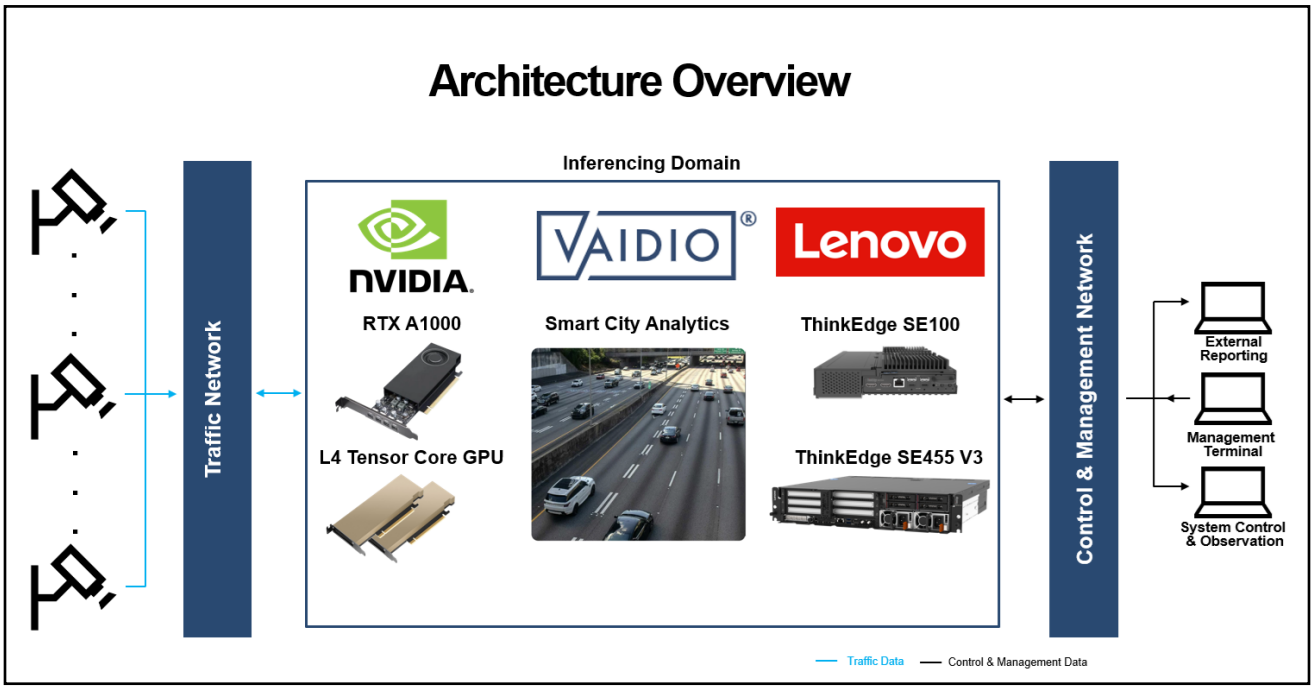

Figure 1. Vaidio AI video analytics on Lenovo servers with NVIDIA GPUs

Video streams are processed through the Vaidio application portal, with AI models accelerated by NVIDIA GPUs on Lenovo SE100 and SE455 V3 servers. GPU acceleration (RTX A1000 and NVIDIA L4) provides the parallel processing power required for running high-density video analytics across entire streaming camera and sensor infrastructure.

This architecture illustrates how the Vaidio AI Vision Platform integrates with Lenovo infrastructure and NVIDIA GPUs to enhance safety and security. Secure connectivity via Port 443 enables integration with the internet, while operators can access analytics and alerts through the Vaidio Portal or mobile devices for real-time monitoring and rapid response.

Figure 2 below illustrates the Vaidio on-premises architecture, showing how video data flows from various sources such as ONVIF IP compliant cameras, the Vaidio CamApp, or third-party NVRs into the Vaidio Core. The system processes video streams through frame processing, detection, analysis, and metadata

generation. Processed metadata and stored video are then made accessible via a web browser or mobile app, enabling users to monitor, search, and analyse recorded and live video in real time.

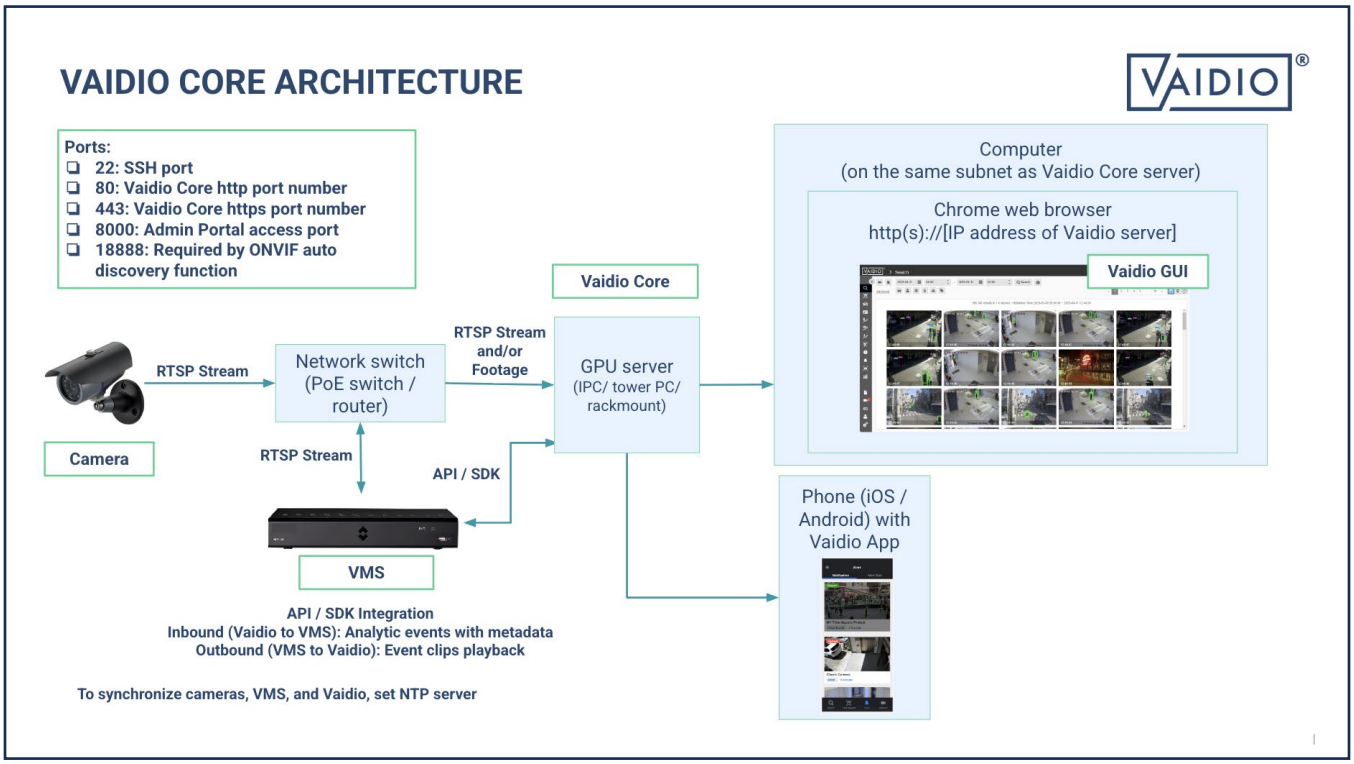

Figure 1. Vaidio on-premises Integration Architecture

Figure 3 shows an architecture overview of a deployment at the edge, built on Lenovo servers with NVIDIA GPU and managed through Kubernetes. It highlights the integration of application with automated software and hardware configuration, supported by network automation and management domain. The architecture ensures simplified deployment and management of applications at the edge while maintaining scalability, security, and centralized control through tools.

With flexible deployment across edge devices and centralized management through Kubernetes, the solution ensures dynamic scaling, secure operations with simplified management from a single node to thousands, in a federated Architecture Component Model.

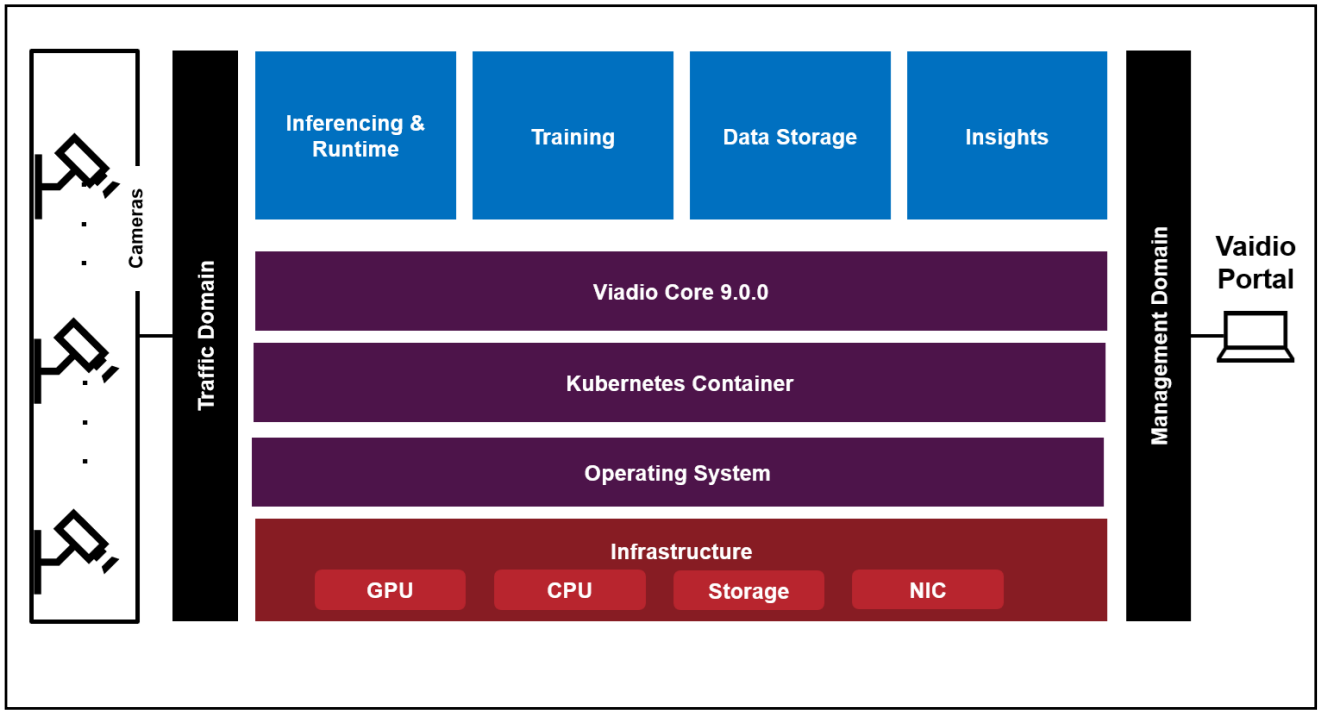

Figure 3. Single-node cluster functional architecture at the edge with Kubernetes

The Vaidio AI Vision Platform is built on a modular architecture that combines advanced video analytics with scalable Lenovo infrastructure. The following key components enable flexible deployment across edge, core, and enterprise environments:

- **Video Sources**
  Includes IP cameras (ONVIF-compliant), the Vaidio CamApp, existing VMS and NVRs platforms, recorded footage and other video feeds that ingest real-time streams into the system.
- **Vaidio Core**
  The central AI engine that performs video decoding, frame processing, object detection, event classification, and metadata generation. Core also manages analytic rules, user authentication, and system configurations.
- **Analytics Modules**
  A library of AI-powered. These applications can be fine-tuned to fit customer use cases.
- **Data & Metadata Storage**
  Stores processed video streams and metadata in scalable databases. Metadata indexing supports rapid search, retrieval, and incident analysis across large video datasets.
- **Management & Control**
  Provides system configuration, policy compliance, alert management, and dashboard visualization. Includes the Vaidio Web Portal and mobile app for operational monitoring and user access.
- **Integration & APIs**
  Open APIs and connectors for integrating with third-party applications such as VMS, PSIM, incident management platforms, and command center systems. Supports RESTful APIs, ONVIF, and standard event forwarding.
- **Infrastructure Layer**
  Powered by Lenovo SE100 and SE455 V2 servers with NVIDIA GPUs for AI acceleration. The

platform is deployed on Kubernetes clusters for elasticity, high availability, and simplified lifecycle management.

Table 1 provides an overview of technologies used in different functional layers as part of this LVD.

Table 1 - LVD Technology Mapping

Layer	Technology	Version	Role
Cameras	IP	Any IP cameras	Provides standardized video streaming and control via ONVIF protocol.
Edge Infrastructure	ThinkEdge	SE100	Optimized for small deployments at the edge.
		SE455 V3	Scalable edge server for medium to large deployments.
GPU	NVIDIA	RTX A1000	Energy efficient GPU for edge AI workloads.
		NVIDIA L4	High-performance GPU for accelerated AI inference and video analytics.
Operating System	Ubuntu 22.04	22.04	Stable, long-term support Linux OS running on edge infrastructure.
CaaS	Kubernetes		Orchestration and lifecycle management of containerized applications.
Application Services	Vaidio Core	9.0.0	Enables Vaidio AI engines for video processing and analytics.
	CUDA	12.9	NVIDIA GPU acceleration for image and video processing.
Dashboard/UI	Vaidio Portal	9.0.0	User interface for managing and controlling Vaidio Core.
Central Management	Vaidio Command Center	9.0	Centralized management for multi-site deployments, scaling across multiple locations.
Statistics/Insights	Vaidio Data	9.0	Aggregates and visualizes metadata from multiple Vaidio nodes; provides monitoring, categorization, dashboards, comparisons, and insights.
Mobile App	Vaidio Mobile App	9.0.0	Allows to receive alerts, perform search and other functions in the mobile app.
Mobile Camera App	Vaidio Cam App	9.0.0	Enables classification and alerts on the mobile phone camera.

Operational Model

The operational model of this LVD is designed to ensure reliable deployment, scalability, and ease of management across diverse environments. It leverages Lenovo servers with NVIDIA GPUs to run Vaidio's AI video analytics platform in a Kubernetes-based architecture, providing flexibility, high availability, and dynamic scaling.

Key Operational Aspects

- **Deployment Flexibility:** Supports deployment across edge devices, centralized servers, or hybrid environments. Kubernetes orchestration allows seamless workload distribution and scaling as video analytics demand fluctuates.
- **High Availability:** Redundancy is built into the architecture with Kubernetes failover capabilities, ensuring minimal downtime and continuous operations even in case of hardware or software failures.
- **Monitoring & Management:** Vaidio provides centralized monitoring dashboards, logs, and health checks for system performance, video processing throughput, and analytics accuracy. Lenovo XClarity integrates with the infrastructure for end-to-end lifecycle management.
- **Security Operations:** Implements secure video streaming, encrypted data storage, and role-based access control for administrators, operators, and viewers. Integration with third-party systems follows ONVIF and RESTful API standards.
- **Integration & Interoperability:** Designed for seamless integration with existing camera networks, VMS/NVRs, and command center applications. APIs and SDKs extend capabilities for custom analytics or external application workflows.
- **Scalability & Growth:** Edge-to-cloud design supports scaling from small deployments with a handful of cameras to city-wide networks with thousands of devices, without compromising performance. LVD further strengthens this by validating different sizing models—small, medium, and large—to guide customer planning and ensure right-sized deployments from pilot projects to full-scale rollouts.

Vaidio Operational Model

Vaidio's video analytics platform integrates video ingestion, real-time analysis, event management, and reporting into a closed-loop workflow.

Core platform functions:

- **Search:** Find and classify vehicles, animals, objects, or items across live streams and files. Filters include time, source, object type, and camera Region of Interest (ROI). Scene details provide contextual understanding with bounding boxes, playback, GPS/indoor maps, and cross-camera tracking.
- **Alerts:** Rules generate real-time notifications for events such as incident detection, idle time, license plate recognition, or person fall. Alerts can be scheduled, throttled with cooldowns, acknowledged, and routed to mobile, VMS, or email.

- **Statistics:** Aggregates detection data into heatmaps, demographics, and business reports, with export options in xlsx, jpg, or pdf.

Mobile apps:

- Vaidio App gives mobile access to search, alerts, scene details, and live view. Supports search, smart tags, and privacy-protected unblur. Real-time alerts are pushed to smartphones.
- Vaidio Cam App turns smartphones into streaming cameras supporting real-time license plate recognition, with mapping IDs to link to Vaidio Core.

Recent enhancements in version 9.0.0:

- AI-based camera auto configuration of analytic engines.
- Expanded cross-camera tracking with GPS and indoor maps.
- Adjustable alert cooldown to prevent alert fatigue.
- UI improvements and expanded VMS/NVR integrations.

Execution workflow:

1. **Ingest video** from IP cameras, files, or third-party systems into the Vaidio Core.
2. **Analyze streams** with video analytics AI agents.
3. **Generate metadata** to enable rapid search and cross-camera tracking.
4. **Trigger alerts** and deliver them to operators or integrated systems via dashboards, mobile apps, or APIs.
5. **Refine analytics** by aggregating long-term data (heatmaps, demographics, reports) and improving accuracy through false detection reporting.

NVIDIA Technology

NVIDIA GPUs are the foundation of Vaidio's real-time AI video analytics, combining 4th-generation Tensor Cores and the energy-efficient Ada Lovelace architecture to accelerate inference at scale while reducing power consumption. The NVIDIA L4 delivers enterprise-class performance with 24GB of GDDR6 memory at just 72 watts, while the RTX A1000 offers 2,304 CUDA cores in a compact, 50-watt design for edge deployments. Built on a 4nm process, the Ada architecture ensures efficient, reliable performance, enabling cities and enterprises to deploy powerful AI analytics at the edge where space, energy, and cost are constrained.

Deployment considerations

Successful deployment of AI-powered video analytics requires careful planning across infrastructure, networking, storage, and system management. Each environment, whether a pilot deployment at the edge or a city-wide rollout, must balance performance, scalability, and security to meet operational needs. LVD with Vaidio streamline this process by providing tested hardware profiles, reference architectures, and configuration guidance that reduce risk and accelerate time-to-value. By aligning compute, storage, and connectivity with the demands of real-time analytics, organizations can ensure consistent performance, simplified management, and long-term scalability across diverse deployment scenarios.

Systems management

The Vaidio Command Center centralizes monitoring, configuration, and updates across deployments from a few sites to thousands of locations. It integrates with IoT devices, VMS platforms, and LDAP, supports 30+ AI analytics models, and enables incident detection with space–time correlation. This unified approach simplifies operations, strengthens security, and ensures scalable, reliable system performance.

Figure 4 shows the capabilities and operational functions of Vaidio Command Center.

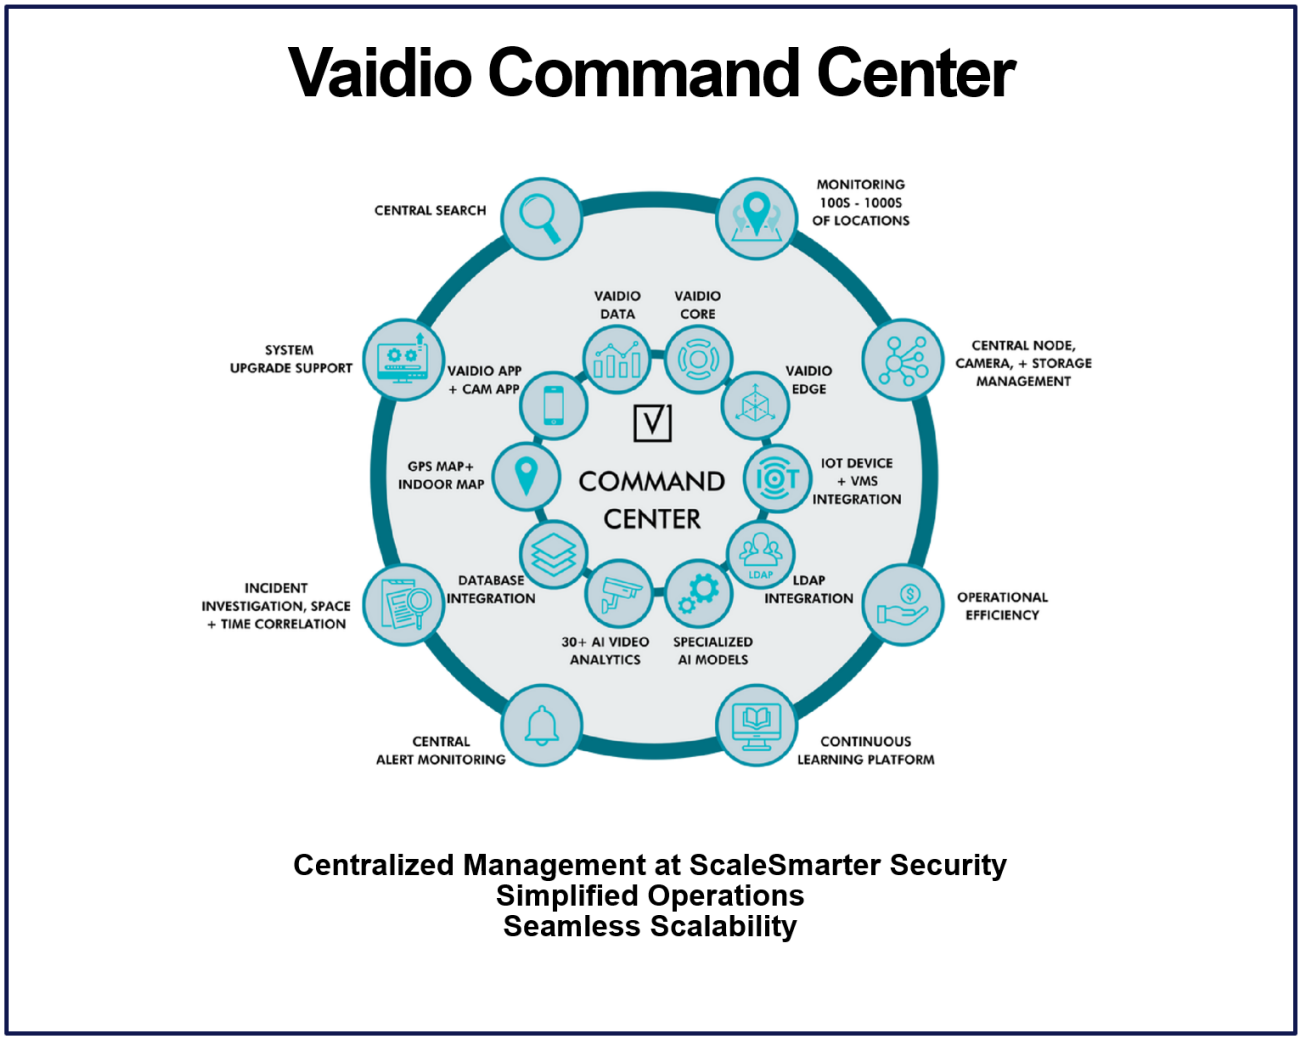

Figure 4 - Vaidio Command Center – Centralized AI Video Analytics and Management Platform

The Vaidio Portal provides a unified interface for monitoring multiple camera feeds with real-time AI analytics. Operators can view live video, detect events such as intrusions or anomalies, and quickly investigate incidents across large-scale environments from a single dashboard. Figure 5 is an example of camera stream processing sampled and classified in Vaidio Portal.

Figure 5 - Multi-Camera Video Monitoring with AI Analytics

Server / Compute Nodes

Lenovo ThinkEdge SE100 Edge Server

The Lenovo ThinkEdge SE100 is a purpose-built server that is one-third the width and significantly shorter than a traditional server, making it ideal for deployment in tight spaces. It can be mounted on a wall, desktop, or rack. The SE100 is optimized for Artificial Intelligence with increased processing power, storage, and networking closer to where data is generated.

Figure 6 visualizes the front-view of Lenovo ThinkEdge SE100 with the GPU Expansion Kit.

Figure 6 – SE100 with attached GPU Expansion Kit

Lenovo ThinkEdge SE455 V3 Edge Server

The Lenovo ThinkEdge SE455 V3 is Lenovo's most powerful edge server, supporting up to two NVIDIA L40 GPUs. Its 2U short-depth chassis can be mounted in either a 2-post or 4-post rack, making it well suited for deployment outside traditional data centers. Designed for demanding AI workloads, the SE455 V3 delivers high-performance video processing and inference at the edge.

Figure 7 visualizes the back-view of Lenovo ThinkEdge SE455 V3.

Figure 7 – SE455 V3

Table 2 shows the main features of the SE100 and SE455 V3.

Feature	SE100	SE455 V3
Form Factor	Ultra-compact (1/3 width, 2.1L base node; wall/desktop/rack-mountable)	2U short-depth rack server (440mm)
GPU Support	1× NVIDIA RTX A1000 (energy-efficient)	Up to 6× single-width GPUs or 2× double-width GPUs (e.g., NVIDIA L4)
Processor	Intel® Ultra 5 class CPU	AMD EPYC™ (up to 64 cores)
Memory	Up to 64 GB	Up to 768 GB
Storage	Encrypted SSD options; small-scale edge storage	Flexible drive bays, >490 TB capacity with Gen5 OCP expansion
Networking	2× 1GbE + management	1/10/25/100 Gb LOM; multiple PCIe NIC options up to 200 Gb
Power	Dual redundant 140W external PSUs (low power design)	Dual redundant 1100W–1800W AC/DC PSUs
Security	TPM 2.0, ThinkShield, intrusion protection	TPM 2.0, ThinkShield, system lockdown, AMD Infinity Guard
Operating Range	5–45 °C, dust/vibration protection, 35 dBA acoustics	5–55 °C, NEBS 3 support, acoustic modes 40/45 dBA
Management	Lenovo XClarity Administrator	Lenovo XClarity Controller / LOC-A
Best Fit	Small edge deployments, space- and power-constrained environments	Medium to large edge AI deployments, GPU-rich video analytics at scale

Networking

Smart city environments must ensure reliable connectivity to support AI-enhanced safety and operational solutions. Robust networking is required to transport RTSP video streams, typically over wired Ethernet or fiber optic infrastructure, and in some cases, high-density Wi-Fi for large public spaces.

Secure connectivity is equally critical, with public IP addressing and integrated firewall protection to safeguard residents, assets, and infrastructure. A practical approach is to leverage existing infrastructure by using port 443 for secure HTTPS communications, enabling outbound encrypted connections for remote support, operations, and ongoing management.

By minimizing disruption, improving operational efficiency, and enabling scalability, this approach ensures that smart cities can effectively deploy and manage AI-powered video solutions across diverse environments.

Storage integration

The SE455 V3 offers flexible storage expansion with optional internal drive bays and an optional second PCIe riser. Its high-speed Gen5 x16 OCP architecture doubles available storage bays, delivering more than 490 TB of total capacity. The platform also provides 20% more PCIe slots compared to competitive edge servers, enabling greater scalability for demanding workloads.

Solution Validation Results

Validation Scope

Testing has been conducted across four categories:

- Object Detection
- Access Control
- Business Intelligence / Statistics / Operations
- Safety and Security / Incidents

Representative use cases included object and vehicle detection, cross-camera tracking, license plate recognition, intrusion and crowd detection, PPE compliance, and smoke/fire monitoring. For a complete list of supported analytics, see Appendix A – Vaidio Analytics Capabilities.

Software Installation

Installing Vaidio Core 9.0 is a straightforward process guided by the Installation Guide, which outlines disk configuration, software dependencies, and network settings. Driver installation is handled by the included “preinstall” script, ensuring all components are correctly set up. As Vaidio Core is containerized, deployment is further simplified with an auto-configuration file that automatically launches both the Vaidio Core service and the Vaidio Web Portal, enabling a smooth and efficient startup.

Testing Methodology

Testing methodology used carefully selected pre-recorded CCTV footage, processed locally on the device to simulate RTSP streams and cover all previously described use case scenarios. For each use case, appropriate video analysis settings were applied to accurately replicate a realistic deployment environment. System resources were continuously monitored while the footage was looped and analyzed by Vaidio Core 9.0. Initial results indicate that Intel Ultra 5 and AMD EPYC processors operated at high utilization, while the NVIDIA L4 and RTX 1000 GPUs maintained moderate utilization, leaving headroom to support additional use cases.

Validation Results

Table 3 – Validation Results

Lenovo Edge Servers	Number of cameras	Use Cases (See Appendix A)
SE100 RTX1000	16 cameras	17 Use Cases
SE455 V3 1xL4	30 cameras	17 Use Cases
SE455 V3 2xL4	70 cameras	17 Use Cases

Vaidio offers 17 specific analytical functions, i.e. use cases, which each camera can perform. (See Appendix A for more info) We observed that the specific use cases applied to each camera have a direct impact on the number of cameras the system can support. The results presented reflect a balanced mix of multiple use

cases running simultaneously. The total number of supported cameras may increase or decrease depending on the combination of use cases selected.

Resource-intensive use cases, such as license plate detection and object tracking, demand more computational power and therefore reduce the maximum number of supported cameras. In contrast, lighter use cases, such as object counting, consume fewer resources and allow for a higher number of cameras to be utilized. These results demonstrate that NVIDIA GPU acceleration directly impacts system scalability, supporting more cameras and use cases per server with consistent responsiveness.

Solution Summary

The Lenovo Validated Design (LVD) with Vaidio provides a proven blueprint for modernizing video infrastructure with AI-driven intelligence. Unlike legacy systems that rely on manual monitoring, this solution automates detection, search, and reporting, enabling cities and enterprises to respond faster, operate more efficiently, and scale with confidence.

Validated on Lenovo ThinkEdge SE100 and SE455 V3 servers with NVIDIA GPU acceleration, the solution delivers real-time analytics across thousands of video streams with sub-second search and investigation capabilities. Its Kubernetes-based architecture ensures seamless integration with existing infrastructure, while supporting flexible deployment from edge pilots to city-wide rollouts.

By combining Lenovo’s high-performance infrastructure with Vaidio’s 30+ advanced analytics, organizations can enhance public safety, improve traffic flow, reduce operational costs, and unlock long-term sustainability in their smart city and enterprise initiatives.

Appendix A: Vaidio Analytics Capabilities

Category	Analytics Functions	Use Cases
People	People counting, crowd detection, loitering, fall detection, cross-camera tracking	Public safety, crowd management, pedestrian safety, law enforcement
Vehicles	Vehicle detection & classification (cars, trucks, buses, motorcycles, bicycles, trains), license plate recognition (LPR), wrong-way detection, congestion analysis	Traffic management, parking optimization
Objects	Object left behind, object removal, scene change detection	Transit hubs, public areas, critical infrastructure
Anomalies & Events	Intrusion detection, suspicious behaviour, PPE compliance, smoke/fire detection	Workplace safety, industrial sites, smart cities
Access	Access control integration	Secure facilities, enterprise campuses
Business Intelligence	Heatmaps, dwell time, smart tags, cross-camera statistics	Retail analytics, urban planning, event management

Appendix B: Lenovo Bill of materials (BOM)

ThinkEdge SE100

Part Number	Product Description	Total Qty
7DGRCTO1WW	**SE100 Small : ThinkEdge SE100 - 3 Year Warranty**	**1**
C31D	ThinkEdge SE100 Chassis	1
BFYE	Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
C30L	ThinkEdge SE100 Planar with Intel Core Ultra 5 225H ,14C, 28W, 1.7GHz	1
C39L	ThinkEdge 16GB TruDDR5 5600MHz (1Rx8) SODIMM	2
C8V3	ThinkSystem M.2 7450 PRO 480GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	1
BS2P	ThinkSystem M.2 7450 PRO 480GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	1
C39N	ThinkSystem NVIDIA RTX A1000 8GB PCIe Gen4 Active GPU	1
C39R	ThinkEdge 140W 230V/115V External Power Supply	1
A4VP	1.0m, 10A/100-250V, C13 to C14 Jumper Cord	1
B755	Desktop Mode	1
C31J	ThinkEdge SE100 Bottom Rubber Feet	1
C31A	ThinkEdge SE100 Fan Module	1
C31C	ThinkEdge SE100 Port Dust Cover Kit	1
BYBQ	XClarity Controller Managed	1
B0MK	Enable TPM 2.0	1
B7XZ	Disable IPMI-over-LAN	1
BB98	Disable IPMI-over-KCS	1
A2N7	Planar Not Integrated With Chassis	1
B0ML	Feature Enable TPM on MB	1
BRPJ	XCC Platinum	1
C3GN	ThinkEdge SE100 LPK	1
BE0B	Non-Redundant	1
C31P	ThinkEdge SE100 Node SSL_LI	1

C31M	ThinkEdge SE100 Node Label GBM	1
C31K	ThinkEdge SE100 Node WW Packaging	1
C8U9	Top-Cover Thermal Gap Pad Kit	1
C319	ThinkEdge SE100 Node Cosmetic Cover	1
C308	ThinkEdge SE100 M.2 Holder	1
C305	ThinkEdge SE100 REGID	1
C316	ThinkEdge SE100 Fan Cable (Bridge Cable)	2
C31T	ThinkEdge SE100 Agency Label	1
C8UC	Front I/O Panel	1
C8UB	Expansion Kit Rubber Feet	1
C8UA	Bottom-Cover Thermal Gap Pad Kit	1
	Auto-Derived Part Items	
C30F	SE100 Expansion Kit for Active Cooling GPU	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

ThinkEdge SE455 V3

Part Number	Product Description	Total Qty
7DBYCTOAWW	**SE455 V3 Medium with 1xL4, Large with 2xL4: ThinkEdge SE455 V3 - 3Yr Warranty for AI**	**1**
BVTK	ThinkEdge SE455 V3 Chassis	1
BFYA	Operating mode selection for: "Maximum Efficiency Mode"	1
BY8K	Acoustic Mode Disabled	1
BY8S	System Operational Temperature 5C to 25C / 41F to 77F	1
BY8X	ThinkEdge SE455 V3 AMD EPYC 8324PN 32C 130W 2.05GHz Processor	1
BW3L	ThinkEdge SE455 V3 1U Heatsink	1
BUVV	ThinkSystem 96GB TruDDR5 4800MHz (2Rx4) 10x4 RDIMM-A	6
5977	Select Storage devices - no configured RAID required	1
C18M	ThinkSystem 2.5" U.2 VA 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	1

BVUU	ThinkEdge SE455 V3 2.5" NVMe Backplane	1
BVUY	ThinkEdge SE455 V3 M.2 SATA/x4 NVMe Adapter with Carrier	1
BXMF	ThinkSystem M.2 PM9A3 3.84TB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	1
BYFH	ThinkSystem NVIDIA L40S 48GB PCIe Gen4 Passive GPU	2
BVUR	ThinkEdge SE455 V3 Riser1	1
BVUS	ThinkEdge SE455 V3 Riser2	1
BMH9	ThinkEdge 1800W 230V Platinum Hot-Swap Gen2 Power Supply	2
BMH2	ThinkEdge 600mm Ball Bearing Rail Kit	1
BS4E	ThinkEdge 130mm USB-C to VGA Display Cable	1
BYBQ	XClarity Controller Managed	1
C1QS	SED Disabled	1
BPKR	TPM 2.0	1
B7XZ	Disable IPMI-over-LAN	1
BB98	Disable IPMI-over-KCS	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
BK15	High voltage (200V+)	1
BVUA	ThinkEdge SE455 V3 Language Pack	1
BVUZ	ThinkEdge SE455 V3 M.2 SATA/x4 NVMe Adapter Onboard NVMe Cable	1
BW1T	ThinkEdge SE455 V3 Riser2 Power Cable	1
BW1S	ThinkEdge SE455 V3 Riser2 PCIe Gen4 Cable	1
BW1R	ThinkEdge SE455 V3 Riser2 PCIe Gen5 Cable	1
BW1Q	ThinkEdge SE455 V3 Riser1 PCIe Gen4 Cable	1
C8DD	ThinkEdge NVIDIA L40S GPU Power Cable Kit	2
BVVW	ThinkEdge SE455 V3 Front Backplane Power Cable	1
BW1P	ThinkEdge SE455 V3 Riser1 PCIe Gen5 Cable	1
BVV5	ThinkEdge SE455 V3 2nd Motherboard Power Cable	1

BVV4	ThinkEdge SE455 V3 1st Motherboard Power Cable	1
BVV3	ThinkEdge SE455 V3 Fan Board Power Cable	1
BVV2	ThinkEdge SE455 V3 Fan Board Sideband Cable	1
BVVP	ThinkEdge SE455 V3 Onboard NVMe Cable for Front Drive Bays	1
BVV0	ThinkEdge SE455 V3 M.2 SATA/x4 NVMe Adapter Power Cable	1
BE0F	N+N Redundancy Without Over-Subscription	1
BW39	ThinkEdge SE455 V3 CPU Air Baffle for 1U Heatsink	1
BW38	ThinkEdge SE455 V3 Supercap Holder	1
BW37	ThinkEdge SE455 V3 M.2 Air Baffle Extension	1
BW36	ThinkEdge SE455 V3 M.2 Air Baffle	1
BVU8	ThinkEdge SE455 V3 LI Service Label	1
BVVY	ThinkEdge SE455 V3 2.5" Drive Bay Filler	3
BY8T	ThinkEdge SE455 V3 OCP Filler	1
BVU4	ThinkEdge SE455 V3 Label Group	1
BVU3	ThinkEdge SE455 V3 System Package	1
BVTP	ThinkEdge SE455 V3 Fan	5
BW3N	ThinkEdge SE455 V3 Security Activation Label	1
BVV1	ThinkEdge SE455 V3 Fan Board	1
BVTM	ThinkEdge SE455 V3 Root of Trust	1
BVUL	ThinkEdge SE455 V3 Power Distribution Board Sideband Cable	1
BVUK	ThinkEdge SE455 V3 Power Distribution Board	1
BW3G	ThinkEdge SE455 V3 1800W Platinum PSU Rating Label	1
BVUJ	ThinkEdge SE455 V3 Regulatory ID	1
BVVJ	ThinkEdge SE455 V3 Riser2 Rear Support	1
BVUH	ThinkEdge SE455 V3 Dust Cover Kit for I/O Ports	1
BVVH	ThinkEdge SE455 V3 Riser1 Rear Support	1
BVVF	ThinkEdge SE455 V3 Riser Side Support	2
BZBL	ThinkEdge SE455 V3 PSU Agency Label	1
C8YD	SE455 V3 Laser Service Indicator	1
	Auto-Derived Part Items	

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

Appendix C: Abbreviations

AI	Artificial Intelligence
API	Application Programming Interface
CaaS	Container as a Service
CCTV	Closed Circuit Television
CPU	Central Processing Unit
CUDA	Compute Unified Device Architecture
GPS	Global Positioning System
GPU	Graphical Processing Unit
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IoT	Internet of Things
IP	Internet Protocol
IT	Information Systems
LDAP	Lightweight Directory Access Protocol
LVD	Lenovo Validated Design
NVR	Network Video Recorder
ONVIF	Open Network Video Interface Forum
OS	Operating System
PPE	Personal Protective Equipment
PSIM	Physical Security Information Management
ROI	Region of Interest
RSTP	Real-Time Streaming Protocol
SDK	Software Development Kit
VMS	Video Management System

Resources

Resources	Links
Vaidio	Vaidio Website
SE100	Lenovo ThinkEdge SE100 Server Product Guide > Lenovo Press
SE455 V3	Lenovo ThinkEdge SE455 V3 Server Product Guide > Lenovo Press
Lenovo XClarity	Systems Management
LOC-A	Lenovo Open Cloud Automation (LOC-A)

Document history

Version 1.0	September 2025	First version includes Lenovo SE100, SE455 V3, and Vaidio 9.0

Trademarks and special notices

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®
ThinkShield®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

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

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

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

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 serves 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 are at your own risk