



Lenovo Validated Design: Public Sector Safety and Security Edge AI Solution

Last update: 10 July 2024

Version 1.0

Provides a technical
overview of Lenovo
ThinkEdge servers

Contains performance data
and sizing recommendations

Shows variety of Kognition
AI Safety and Security use
cases

Explains the benefits of Edge AI
solutions for enhancing public
sector safety and security

Vanita Meyer
Chris Arenas



Table of Contents

1. Introduction.....	3
2. Key Challenges and Business Value	4
2.1 Key Challenges.....	4
2.2 Business value.....	5
2.3 Benefits of the Solution	6
3. Technology Overview	7
3.1 Functional requirements	7
3.1.1 Hardware Requirements	7
3.1.2 IP Cameras.....	9
3.1.3 Qualcomm AI 100 Accelerator	9
3.1.4 VMware	10
3.2 Non-functional requirements.....	10
4. Design Overview	12
4.1 Lenovo Validated Design Architecture	12
4.2 Kognition AI Architecture	12
5. Deployment Model	15
6. Deployment considerations.....	16
6.1 Systems Management	16
6.2 Server/Compute Nodes	16
6.3 Networking.....	16
6.4 Storage Integration	16
6.5 Sizing Guide.....	17
6.6 Importance of Sizing Recommendations	17

6.7	Compute Resource Requirements.....	18
6.8	Collaboration with Lenovo.....	18
7.	Validation Scope.....	19
8.	Test and Validation Methodology and Results.....	21
8.1	Methodology	21
8.2	Validated Results	22
9.	Implementation Guide	23
10.	Life Cycle Management Software	25
10.1	Lenovo Open Cloud Automation (LOC-A).....	25
10.2	Lenovo XClarity.....	26
11.	Appendix:	28
11.1	Hardware components	28
11.2	Lenovo Bill of materials.....	31
11.3	Abbreviations	37
12.	Resources	39
13.	Document history	40
14.	Trademarks and special notices	41

1. Introduction

Edge computing is transforming data processing by handling data closer to its origin, reducing reliance on distant cloud services. This approach is crucial for scenarios requiring immediate analysis or minimal latency. Gartner predicts that by 2025, over 50% of data will be managed outside traditional data centers, boosting the edge computing market. Organizations are strategically placing workloads near data generation points to leverage real-time insights and avoid cloud processing delays.

Revolutionizing Public Sector Safety with AI and Advanced Technologies

AI is enhancing public sector safety through automated data analysis and real-time decision-making. By analyzing data from CCTV, social media, and public records, AI can identify crime patterns and predict incidents, improving efficiency and response times. Public sector organizations are also prioritizing cybersecurity with robust security frameworks and zero-trust models to protect sensitive data.

The adoption of edge computing, computer vision, and IoT sensors in public safety is growing. These technologies, along with mobile access, smart sensors, and automated systems, are enhancing response times and situational awareness.

Overview of Lenovo Edge Computing Solution with Kognition

This document provides an overview of Lenovo's edge computing solution for public sector safety and security, featuring Kognition running on Lenovo ThinkEdge servers SE360 V2, SE450, and SE455 V3. The objective of this Lenovo Validated Design (LVD) is to assess the performance of Kognition on each platform. The results will yield critical data to develop sizing recommendations. The solution is rigorously tested and well-documented to ensure reliability and consistency in customer deployments.

This paper offers planning and design considerations, along with best practices for implementing public sector Kognition Safety and Security solutions using Lenovo products.

Intended Audience

The intended audience for this document includes IT professionals, technical architects, sales engineers, and consultants. It aims to assist in the planning, designing, and implementation of public sector Kognition Safety and Security solutions.

2. Key Challenges and Business Value

2.1 Key Challenges

Public safety and security, particularly in public schools, face several significant challenges. Here are some of the current problems:

- **Gun Violence and School Shootings:** The increasing frequency of gun violence and school shootings is a major concern. Schools are struggling to find effective measures to prevent such incidents and protect students and staff. The need for improved security infrastructure and preventive measures is critical.
- **Mental Health Issues:** There is a growing recognition of the link between mental health and safety. Many students face mental health challenges that can lead to violence or self-harm. Schools often lack the resources and trained personnel to provide adequate mental health support, which can exacerbate safety issues.
- **Cybersecurity Threats:** As schools increasingly rely on digital platforms for learning and administration, they become targets for cyberattacks. These attacks can compromise sensitive student information and disrupt educational activities. Schools need better cybersecurity measures to protect against these threats.
- **Bullying and Harassment:** Bullying, including cyberbullying, continues to be a significant issue in schools. It can lead to severe emotional and psychological harm, affecting students' well-being and academic performance. Effective anti-bullying programs and policies are essential to address this problem.
- **Emergency Preparedness and Response:** Schools often lack effective emergency preparedness plans and response strategies. This includes drills for active threat situations, natural disasters, and other emergencies. Ensuring that staff and students are well-prepared for various scenarios is crucial for minimizing harm during actual events.
- **Communication Barriers:** Effective communication is essential for safety, but many schools face challenges in maintaining clear and efficient communication channels during emergencies. This can include outdated communication systems and lack of coordination with local law enforcement and emergency services.

Sources: see appendix

2.2 Business value

Kognition AI provides solutions through its advanced artificial intelligence platform, which enhances safety, security, and operational efficiency in public sector environments. The platform integrates and automates security cameras, access control systems, and IoT sensors to form a dynamic, real-time alerting and analytics system. This system is designed to elevate security by generating actionable alerts to prevent theft, enhance workplace safety, and improve overall store management.

The use of Kognition's AI-driven technology in public sector settings facilitates several key functions:

- **Access Control:** Regulate access to restricted areas by verifying identities and detecting unauthorized entry attempts.
- **Activity Analysis:** Analyze behavior patterns to detect anomalies or suspicious activities, enabling proactive security measures.
- **Crowd Formation / Flash Mob Alert:** Detect the formation of large crowds or potentially volatile gatherings, alerting security personnel for crowd control efforts.
- **Identity Verification:** Validate identity of authorized individuals using facial recognition.
- **Intrusion Detection:** Unauthorized entry or tampering is detected, providing early warnings to security teams.
- **Object Detection:** AI technology enhances surveillance by identifying and categorizing objects within a scene.
- **Occupancy / People Counting:** Accurately count and track people and visualize with interactive charts.
- **Panic Button / Alarms:** Silent and mobile duress alarms integrated with CCTV and alarm system.
- **Smoke & Fire Detection:** Monitor for signs of smoke or fire, enabling a rapid response to potential emergencies.
- **Threat Detection and Response:** Threats can be detected in real-time and appropriate response protocols can be triggered.
- **Watchlist Monitoring:** Individuals can be monitored against watchlists for known threats or persons of interest, with alerts triggered upon a match.
- **Weapons Detection:** Alert security personnel to the presence of weapons in monitored areas.
- **Vehicle & License Plate Recognition:** Vehicle identification, counting and parking lot automation.

2.3 Benefits of the Solution

Lenovo Edge computing and Kognition AI solutions provide numerous benefits for enhancing the Kognition Safety and Security of the public sector. Key benefits include:

- **Lenovo Edge Computing platforms:** These are powerful to support the Kognition AI solution, cost-effective, and energy-efficient, ensuring it serves as a long-term solution.
- **Integration with Existing Infrastructure:** Kognition platform works seamlessly with existing cameras, sensors, and access control systems, eliminating the need for costly replacements. This makes it easy to upgrade to a smart building with minimal disruption and expense.
- **Real-Time Monitoring and Alerts:** The system provides real-time video analytics and object detection, generating actionable alerts for various threats such as unauthorized access, intrusions, and safety hazards. This proactive approach helps prevent incidents before they escalate.
- **Enhanced Workplace Safety:** Kognition enhances workplace safety by monitoring for unsafe conditions and alerting staff to potential hazards. This includes threat detection, detection of slip and fall incidents, perimeter breaches, and other safety violations, helping to protect workers and visitors.
- **Scalability and Flexibility:** Kognition and Lenovo Edge solutions are highly scalable, capable of supporting thousands of cameras and sensors across large facilities. This scalability ensures that the system can grow with the needs of the organization.
- **Privacy and Cybersecurity:** The platform includes robust privacy and cybersecurity features, such as data encryption, privacy masking, and customizable data retention policies. This ensures that sensitive data is protected, and that the system complies with privacy regulations.

By leveraging these benefits, organizations can significantly enhance their security posture, improve safety outcomes, and ensure compliance with regulatory requirements, all while maintaining flexibility and scalability to meet future needs.

3. Technology Overview

This section provides a comprehensive exploration of the technologies deployed and the requirements addressed in public sector Kognition Safety and Security solutions.

3.1 Functional requirements

Kognition AI generates real-time, actionable alerts to prevent espionage, hacking, theft, disease spread, active threat scenarios, and other significant threats. It proactively enhances building Kognition Safety and Security while offering occupants enriched smart digital experiences. Leveraging AI-powered cognitive services, Kognition analyses data from sensors and cameras to detect patterns, forecast crowd movements, and identify potential bottlenecks or safety issues. This proactive data processing allows authorities to promptly address risks, manage crowds, and bolster overall security and safety for employees, students, and visitors.

For testing we utilized several of Kognition's use cases to thoroughly test and monitor how both our platform and their software handled peak loads. The use cases utilized include facial recognition, threat detection, behavioral analytics, object identification and watchlist. Since the system is scalable our testing included 1-100 concurrent streams, the maximum of which changed based on sizing. The purpose was to demonstrate several factors, but most notably:

- Faster response time and alerting to potential threats
- Efficient use of system resources
- System preparedness and resiliency

3.1.1 Hardware Requirements

Table 1: ThinkEdge SE360 V2 Hardware Requirements

Requirements for Validation	1080p IP cameras (5 -10 streams)
	VMware
ThinkEdge SE360 V2 Deployment	16 Core Intel Xeon-D Processor (Hyper-thread Capable)
	128 GB of DDR4 3200 MHz RAM
	2 TB of NVMe Storage

	Qualcomm Cloud AI 100 Accelerator
	Nvidia A2 Accelerator
	4x1 GB Ethernet
Interoperability	This validation is based on hosting Kognition AI applications on this platform. Clients can consolidate additional applications by virtualizing the environment. In such cases, an interoperability test is required to ensure that all ISV applications can operate seamlessly on the same platform.
Performance Expectations/Threshold	Expected utilization of 75% of shared resources (i.e - CPU, RAM, HD)

Table 2: ThinkEdge SE450 Hardware Requirments

Requirements for Validation	1080p IP cameras (0 – 50 streams)
	VMware
ThinkEdge SE450 Deployment	32 Core Intel Xeon Platinum 8352Y (Hyper-thread Capable)
	128 GB of DDR4 3200 MHz RAM
	6 TB of NVMe Storage
	Qualcomm Cloud AI 100 Accelerator x 2
	4x1 GB Ethernet
Interoperability	This validation is based on hosting Kognition AI applications on this platform. Clients can consolidate additional applications by virtualizing the environment. In such cases, an interoperability test is required to ensure that all ISV applications can operate seamlessly on the same platform.
Performance Expectations/Threshold	Expected utilization of 75% of shared resources (i.e - CPU, RAM, HD)

Table 3: ThinkEdge SE455 V3 Hardware Requirements

Requirements for Validation	1080p IP cameras (0 -100 streams)
ThinkEdge SE455 V3 Deployment	64 Core AMD EPYC 8004 (Hyper-thread Capable)
	256 GB of DDR4 3200 MHz RAM
	6 TB of NVMe Storage
	Qualcomm Cloud AI 100 Accelerator x 2
	4x1 GB Ethernet
Interoperability	This validation is based on hosting Kognition AI applications on this platform. Clients can consolidate additional applications by virtualizing the environment. In such cases, an interoperability test is required to ensure that all ISV applications can operate seamlessly on the same platform.
Performance Expectations/Threshold	Expected utilization of 75% of shared resources (i.e - CPU, RAM, HD)

3.1.2 IP Cameras

The Kognition Safety and Security Edge solution leverages existing IP camera infrastructure within the spaces to enhance safety, security, operations, and retail analytics and many other use cases. Using existing cameras for edge AI involves leveraging the current surveillance or operational cameras within a facility to perform AI-driven analytics directly at the site of data collection, without needing to transmit large amounts of data back to a central server or cloud. The cameras utilized for testing have a resolution of 1080p, ensuring high-quality video capture. Additionally, the application ingests video at a rate of 10-15 frames per second (fps), which balances the need for real-time processing and data efficiency. This setup is particularly effective for real-time processing needs where low latency and immediate action are crucial.

3.1.3 Qualcomm AI 100 Accelerator

The Kognition Safety and Security solution leverages the Qualcomm Accelerator on the Lenovo ThinkEdge SE360 V2 server enabling faster on-site processing of video data captured by the store cameras. It reduces

latency, minimizes bandwidth consumption, and potentially enhances security, empowering Kognition to deliver faster and more effective functionalities for the Edge solution. The Qualcomm Accelerator acts like a powerful engine within the server, specifically designed for tasks like video analysis using AI.

3.1.4 VMware

VMware provides virtualization and cloud infrastructure solutions that help businesses optimize their IT environments. With VMware's suite of products, organizations can create virtual machines (VMs) that run multiple operating systems and applications on a single physical server, improving resource utilization and reducing costs. VMware's capabilities include cloud management, enabling integration with public and private clouds, and offering tools for automation, orchestration, and security. Additionally, VMware features products like vSphere for efficient server management, NSX for network virtualization, and vSAN for software-defined storage, making it a versatile solution for various IT infrastructure needs.

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. Kognition will alert to detected threat within 3-5 seconds of the threat appearing on the camera.
- **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 3-5Mbps of bandwidth, depending on frame rate and bitrate. Software specific requirements will be attached to the appendix.

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 responsive 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

4.1 Lenovo Validated Design Architecture

The diagram outlines the IT infrastructure supporting Kognition AI Kognition Safety and Security solutions. The system handles various applications such as watchlist alerts, threat detection, intrusion detection, crowd forming, smoke and fire detection, and objects left behind. The infrastructure operates on Ubuntu 20.04 LTS with Kubernetes/Docker for container management.

The deployment options include both bare metal and virtualized environments, specifically using SE455 V3 and SE360 V2 for bare metal setups, and SE450 for VMware setups. The edge systems involved are SE360 V2, SE450, and SE455 V3, featuring components like CPUs, accelerators, memory, and I/O interfaces to ensure robust and efficient processing capabilities.

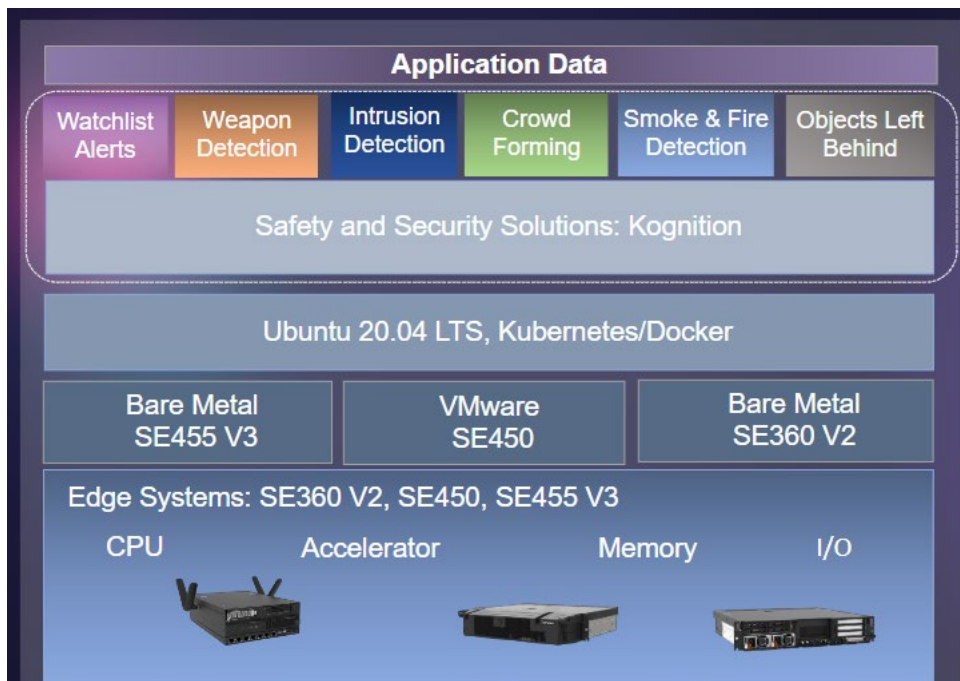


Figure 1: Lenovo Validated Design for Public Sector Kognition Safety and Security Edge AI

4.2 Kognition AI Architecture

The Kognition AI architectural diagrams provide a comprehensive visual representation of the system's design and infrastructure. These diagrams are essential for understanding how the various components of Kognition AI interact and integrate within your business environment. They detail the flow of data, the relationship between different modules, and the overall structure of the application, offering a clear roadmap for deployment and scalability.

Key Features of the Architectural Diagrams:

- **Component Interactions:** The diagrams illustrate how Kognition AI's components, such as data ingestion, processing, machine learning models, and output generation, work together seamlessly. This helps in identifying key touchpoints and integration points within your existing systems.
- **Data Flow:** Clear visualization of data pathways ensures that users understand how data moves through the system, from initial capture to final analysis and reporting. This is crucial for optimizing performance and identifying potential bottlenecks.
- **Infrastructure Layout:** The diagrams provide an overview of the required infrastructure, including compute resources, storage solutions, and network configurations. This aids in planning and deploying the necessary hardware and software components.

KOGNITION ARCHITECTURE

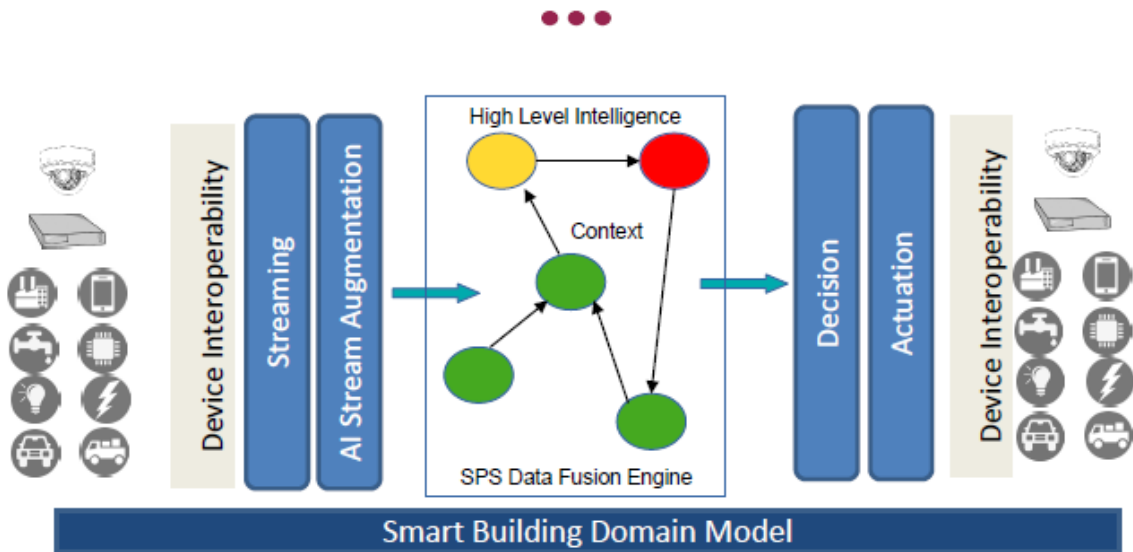


Figure 2: Kognition Architecture

Figure 3 below illustrates the decision-making process for the Kognition application. Real-time data is fed into the system, where robust and innovative AI inferencing rapidly processes the information on-premises.

Subsequently, the application triggers alerts based on predefined rules set on the dashboard and disseminates these alerts to the relevant personnel.

Unified Physical Security Operations & Intelligence Platform

Platform is architected as a software-defined framework for physical security intelligence which is unmatched by any competitor

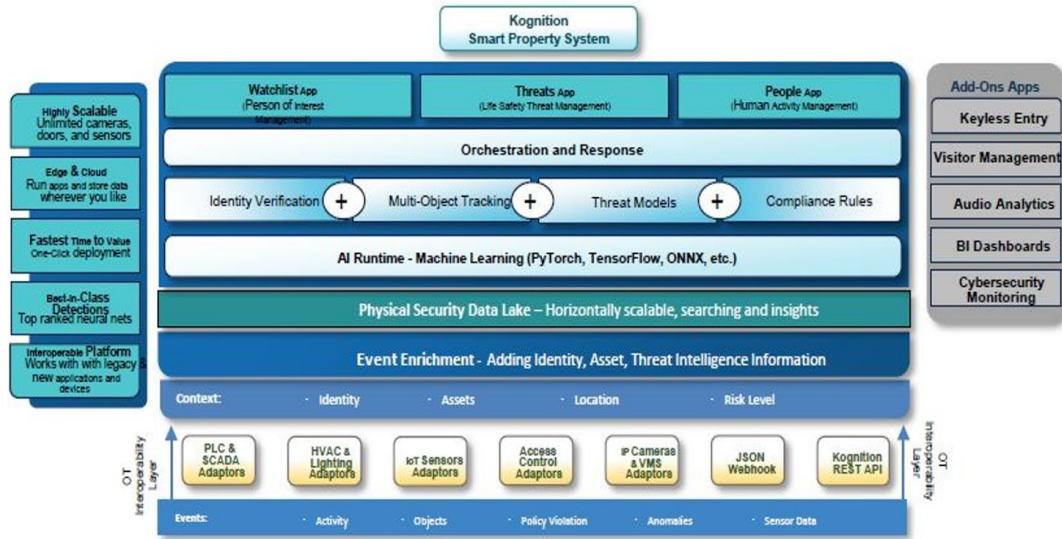


Figure 3: Component Relationship Diagram

5. Deployment Model

The Kognition AI deployment model offers a robust, scalable, and flexible solution designed to integrate seamlessly into various business environments. Leveraging cutting-edge machine learning algorithms and advanced analytics, Kognition AI provides real-time insights and automation capabilities tailored to enhance operational efficiency and decision-making processes. The deployment model supports both cloud-based and on-premises installations, ensuring data security and compliance with industry standards.

A key feature of the Kognition AI deployment model is its use of containerization and Kubernetes orchestration. By utilizing containerization, Kognition AI ensures consistent and reliable application performance across different environments, facilitating easier scaling and management of services. Kubernetes further enhances this by automating deployment, scaling, and operations of application containers, providing a resilient and high-availability infrastructure.

With a focus on adaptability, Kognition AI can be customized to meet specific industry needs, from finance and healthcare to manufacturing and retail. Its user-friendly interface and comprehensive support services make the transition to AI-driven operations smooth and efficient, empowering businesses to unlock new levels of productivity and innovation.

For a detailed illustration of the Kognition AI deployment model, including its architecture and workflow, please refer to Figure 4.

Kognition Deployment Architecture

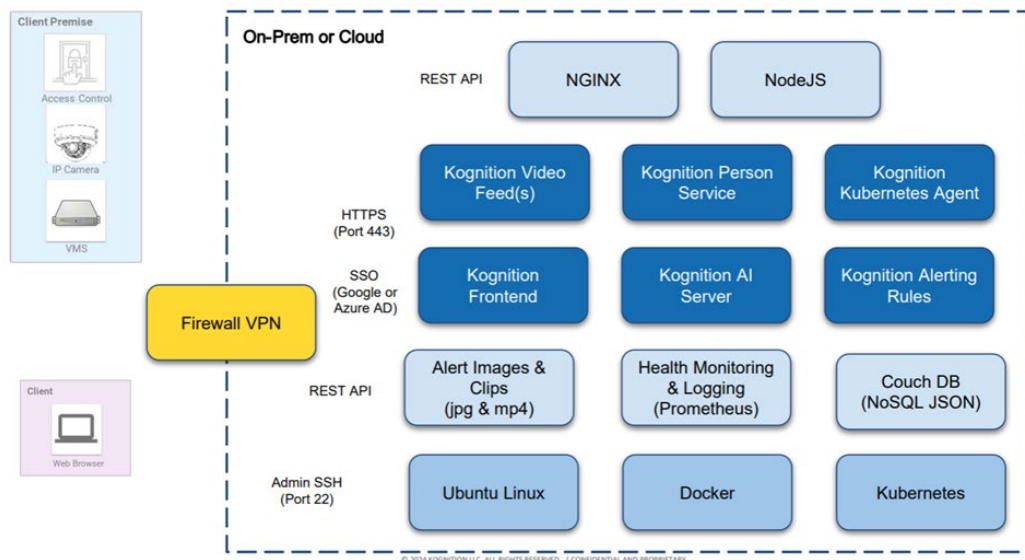


Figure 4: Kognition Deployment Architecture

6. Deployment considerations

Deploying Kognition 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 Kognition AI application. Implement robust monitoring and logging to track system performance and quickly identify issues. Utilize automated tools for updates, patches, and configurations to keep the application secure and up to date. Implement access control and role-based management to protect sensitive data and system integrity. Kognition runs on standard Ubuntu 20.04 LTS and can be managed using standard tools. Kognition also includes Prometheus monitoring and Grafana reports for various system level metrics.

6.2 Server/Compute Nodes

Deploying Kognition 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. Kognition 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 Kognition AI deployment. Use a high-speed network infrastructure to support data transfer between application components. Implement a hybrid network architecture that accommodates both on-premises and cloud-based resources. Enhance network security with firewalls and encryption to protect data integrity and prevent unauthorized access. An average HD camera utilizes 2-5Mbps of bandwidth, depending on frame rate and bitrate. A 4k stream can consume 25Mbps. Depending on the number of video streams, properly sized network switches and network interface cards are required. Kognition recommends a minimum of a 1Gbps network backbone.

6.4 Storage Integration

Efficient storage integration is vital for managing the large volumes of data processed by Kognition 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 Kognition AI application, maximizing its benefits and capabilities.

6.5 Sizing Guide

Kognition and Lenovo have collaborated to develop a comprehensive sizing guide for compute resources, ensuring that your deployment of Kognition AI is optimized for performance and efficiency. Proper sizing of compute resources is crucial to handle the demanding workloads of AI applications and to achieve the best possible outcomes.

Table 4: ThinkEdge SE360 V2, SE450, SE455 V3 Sizing Guide

Model	Cameras	Proposed HW	CPU	RAM	Storage	GPU	Boot	OS	NIC
XS	6	SE360 V2	8C Xeon	64GB	2x 960GB	1x AI 100	2x 240GB	Ubuntu 20.04 LTS	2x 1GbE
S	12	SE360 V2	12C Xeon	128GB	2x 960GB	1x AI 100	2x 240GB	Ubuntu 20.04 LTS	2x 1GbE
S+	25	SE360 V2	16C Xeon	128GB	2x 960GB	1x AI 100	2x 240GB	Ubuntu 20.04 LTS	2x 1GbE
M	50	SE450	32C Xeon	128GB	3x 1.92TB	2x AI 100	2x 480GB	Ubuntu 20.04 LTS	4x 1GbE
M	50	SE455 V3	32C Epyc	128GB	3x 1.92TB	2x AI 100	2x 480GB	Ubuntu 20.04 LTS	4x 1GbE
L	100	SE455 V3	1x 64C Epyc	256GB	3x 1.92TB	2x AI 100	2x 480GB	Ubuntu 20.04 LTS	4x 1GbE

6.6 Importance of Sizing Recommendations

Accurate sizing recommendations are essential to ensure that your infrastructure can support the intensive data processing and AI inferencing tasks required by Kognition AI. By following the guidelines developed by Kognition and Lenovo, businesses can avoid common pitfalls such as under-provisioning, which can lead to performance bottlenecks, and over-provisioning, which can result in unnecessary costs. The sizing recommendations, outlined below, should be used as a guideline for implementation. Each customer's actual deployments can vary based on specific requirements, in such cases sizing can be modified to accommodate those requirements.

6.7 Compute Resource Requirements

Kognition AI applications require a robust and balanced compute environment. The sizing guide outlines specific requirements for CPUs, GPUs, and memory to handle various AI workloads efficiently. High-performance CPUs are needed for general data processing, while powerful GPUs are critical for training and inferencing machine learning models. Adequate memory is also necessary to ensure smooth operation and quick access to data.

6.8 Collaboration with Lenovo

The partnership between Kognition and Lenovo leverages Lenovo's expertise in hardware and infrastructure solutions to provide tailored recommendations for Kognition AI deployments. This collaboration ensures that the sizing guide is based on real-world testing and industry best practices, offering reliable and validated guidelines for businesses.

By adhering to the compute resource sizing recommendations from Kognition and Lenovo, businesses can achieve optimal performance, scalability, and cost-efficiency in their AI deployments, ensuring that they fully leverage the capabilities of Kognition AI.

7. Validation Scope

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

- **Speed and Reliability:** The designs are engineered to allow rapid deployment of application-optimized edge. This speeds up the process by reducing the research, design, and troubleshooting time typically required for setting up such infrastructures.
- **Customized Solutions:** The architecture is designed to be flexible, allowing adaptations based on specific scenarios. This customization capability ensures that each deployment can meet the precise needs of different environments, enhancing both efficiency and effectiveness.
- **Comprehensive Testing:** This design has gone through testing and validation by Lenovo resources, in collaboration with third-party software and hardware partners. This rigorous testing helps reduce the customer's proof-of-concept time and eliminates the lengthy periods usually required to design and test configurations.

The validation scope that involves Lenovo hardware and Kognition application focuses on ensuring that this solution meets specific performance and interoperability standards suitable for various deployment scenarios. Here's a breakdown of what this involves:

Performance Testing:

Performance testing in this context aims to assess how well the Lenovo hardware can run software applications provided by Kognition under typical or extreme conditions. This testing usually includes:

- **Load Testing:** Evaluating the system performance under expected and peak load conditions to ensure the hardware and software can handle anticipated user actions without performance degradation.
- **Stress Testing:** Pushing the system beyond normal operational capacities to see how it handles extreme workload, testing the limits of system components in terms of robust handling and error management.

These tests are critical to determine the practical capabilities of the hardware when running complex Kognition applications, ensuring that the solutions can support the necessary business processes effectively and efficiently.

Interoperability Testing:

Interoperability testing involves verifying that Lenovo hardware effectively integrates with software from Kognition and that all components of the system work harmoniously together. Key aspects include:

- **Compatibility Testing:** Ensuring that the software operates as expected on the Lenovo hardware across various systems and environments without causing conflicts or errors.
- **Integration Testing:** Checking the connections between the Kognition applications and other system components to ensure data flows seamlessly across the platform, confirming that APIs, data formats, and protocols are correctly aligned.
- **Network Interoperability:** Testing the performance and reliability of the network connections between the software applications on Lenovo hardware, particularly important for cloud-based or distributed systems.

This testing is crucial to ensure that when a business implements a solution involving Lenovo hardware and Kognition applications, the components will function correctly together, providing a stable and reliable system.

These tests are part of a broader validation process that aims to deliver high assurance to businesses that their investment in technology will yield a robust, efficient, and compatible IT environment, ready to meet their operational demands.

8. Test and Validation Methodology and Results

The Kognition Safety and Security Edge solution was validated in 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 at 10-15 fps. Optimized network conditions ensured efficient AI-driven analytics. Use cases included person identification, threat detection, and object detection.

Different use cases can impact the number of concurrent camera streams the system can handle. Variables such as camera types, network bandwidth, and deployment models at a customer's site may affect performance, influencing the system's effectiveness and responsiveness in real-world deployments.

8.1 Methodology

The testing methodology for the Kognition.ai application was designed to rigorously evaluate its capabilities in object detection, threat detection, and facial recognition. We meticulously validated the application on various servers, including the SE360 V2, SE450, and SE455 V3, 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.

Our object detection tests assessed the application's ability to accurately identify and categorize objects within a video stream. For threat detection, we evaluated the application's proficiency in recognizing and flagging potentially dangerous situations or behaviors. The facial recognition tests focused on the application's accuracy and reliability in identifying and verifying individual faces.

By leveraging the SE360 V2, SE450, and SE455 V3 servers, we tested the Kognition.ai application's performance under various conditions and configurations. This thorough testing process validated the application's performance, reliability, and scalability.

8.2 Validated Results

Table 5 below presents the validated results from our rigorous testing of the Kognition.ai application, as previously discussed. These results were obtained by evaluating the application’s performance in object detection, threat detection, and facial recognition, using the SE360 V2, SE450, and SE455 V3 servers. For more detailed results, please contact your sales representative.

Table 5: Test and validation results

Sizing Model	Platform	Cameras	CPU Utilization		Storage IOPS	RAM Utilization		Storage Utilization	Network Utilization	Dashboard Response	Maximum Camera Streams
			Normal	Peak		Normal	Peak				
XS	SE360 V2	6	Please contact your sales engineers and sales reps for the in-depth data.								
S	SE360 V2	12									
S+	SE360 V2	25									
M	SE450	50									
M	SE455 V3	50									
L	SE455 V3	100									

9. Implementation Guide

Deploying Kognition AI on Lenovo ThinkEdge V2, ThinkEdge SE450, and ThinkEdge SE455 V3 servers involves specific prerequisites and steps to ensure a smooth installation process. Each server model requires Ubuntu 20.04 LTS (Focal Fossa) as the base operating system, with the ThinkEdge SE450 also utilizing VMware for virtualization. Additionally, all three systems are equipped with Qualcomm AI 100 accelerators to enhance AI performance.

Lenovo ThinkEdge SE360 V2 and SE455 V3

1. Install Hardware Components:

- **RAM:** Install 128 GB of RAM on the Lenovo ThinkEdge SE360 V2.
 - **Storage:** Install 2 TB of NVMe storage on the SE360 V2 and 6 TB of storage on the SE455 V3.
 - **Qualcomm AI 100 Accelerators:** Ensure that the Qualcomm AI 100 accelerators are properly installed and configured to optimize AI processing capabilities.
2. **Prepare the Server:** Ensure that the Lenovo ThinkEdge SE360 V2 or SE455 V3 server is equipped with Ubuntu 20.04 LTS (Focal Fossa). This version provides the stability and support needed for running Kognition AI.
 3. **Remote Installation:** The Kognition implementation team remotely connects to the servers via SSH to handle the installation. They will download and install Kognition AI directly onto the server, ensuring all necessary configurations are applied.
 4. **Verify Installation:** After installation, the implementation team will verify that Kognition AI is running correctly and perform any necessary initial setup or configuration.

Lenovo ThinkEdge SE450

1. Install Hardware Components:

- **RAM:** Install 128 GB of RAM on the Lenovo ThinkEdge SE450.
 - **Storage:** Install 6 TB of storage to handle large datasets and AI models.
 - **Qualcomm AI 100 Accelerators:** Ensure that the Qualcomm AI 100 accelerators are properly installed and configured to optimize AI processing capabilities.
2. **VMware Setup:** Install VMware on the Lenovo ThinkEdge SE450 server to enable virtualization. Configure VMware to host the necessary virtual machines (VMs) for Kognition AI.
 3. **Prepare the Server:** Ensure that the VMs are running Ubuntu 20.04 LTS (Focal Fossa) as the base operating system.

4. **Remote Installation:** Within the VMs, the Kognition implementation team will remotely connect via SSH to download and install Kognition AI, ensuring all necessary configurations are applied.
5. **Verify Installation:** After installation, the implementation team will verify that Kognition AI is running correctly on the VMs and perform any necessary initial setup or configuration.

By ensuring these hardware components are installed first and utilizing the expertise of the Kognition implementation team for remote installation, businesses can effectively deploy Kognition AI on Lenovo ThinkEdge SE360 V2, ThinkEdge SE450, and ThinkEdge SE455 V3 servers. Leveraging the robust capabilities of each hardware platform, including Qualcomm AI 100 accelerators and ample storage, will maximize the performance and efficiency of your AI operations.

10. Life Cycle Management Software

By utilizing Lenovo Open Cloud Automation and Lenovo XClarity, organizations can achieve more efficient, reliable, and scalable management of their cloud and Edge infrastructures, ultimately leading to improved operational efficiency and reduced costs.

10.1 Lenovo Open Cloud Automation (LOC-A)

Lenovo Open Cloud Automation is a no-code software solution that helps to securely accelerate the deployment of distributed Edge infrastructure, at scale. With Lenovo Open Cloud Automation, distributed Edge infrastructure is deployed faster, with fewer resources required while also reducing the overall CO2 emissions of the process.

Lenovo Open Cloud Automation achieves this by embodying a simplified deployment process and by automating the steps of the process too complex for technicians or engineers to execute without mistake. This delivers a standardized experience for both the infrastructure admin team and the field technicians, while still being customer, deployment, and site specific. Here are its primary functions:

Near-Zero Touch Provisioning and on-boarding of bare-metal Edge nodes:

- Implements a process to enable late-binding of Edge Servers to their owners and automates the provisioning and configuration of the basic
- Day-0 settings of the servers in the field before onboarding them. Once onboarded, the Edge nodes will be used for hosting Edge Cluster software and Edge applications. The base OS is also deployed remotely by LOC-A.

Automated Deployment:

- Automates the installation and configuration of distributed Edge infrastructure, from bare metal to production ready IaaS/CaaS.
- Supports secured, rapid and consistent deployment of distributed Edge infrastructure by applying configuration patterns on all levels of the infrastructure (HW, OS, Cluster).

Scalability:

- Provides tools to efficiently scale the deployment of edge infrastructure, allowing for the addition or removal of Edge nodes as needed.
- Supports deployments from small-scale to very large-scale with a flexible architecture that leverages Kubernetes.

Lifecycle Management:

- Offers comprehensive management of the edge infrastructure lifecycle, including deployment, redeployment, and cluster expansion.
- Ensures consistent and reliable operation of Edge services.

Orchestration and Integration:

- Integrates with various Edge-cloud management platforms and orchestration tools, enabling automated workflows and processes.
- Supports integration with Edge application orchestrators for full-stack enablement and user experience.

LOC-A enables rapid, large-scale infrastructure provisioning without the need for golden images or staging environments. By significantly reducing transport-related carbon emissions and virtually eliminating the need for truck rolls in edge deployments, LOC-A supports IT decarbonization initiatives.

10.2 Lenovo XClarity

Lenovo XClarity™ is a suite of management tools that include centralized resource management solution, aimed at reducing complexity, speeding response, and enhancing the availability of Lenovo® server systems and solutions. Here are its key features:

Centralized Management:

- Provides a single, unified interface for managing Lenovo servers, storage, and networking components.
- Simplifies infrastructure management by consolidating various management tasks into one platform.

Monitoring and Alerts:

- Offers real-time monitoring of hardware health and performance.
- Provides alerts and notifications for potential issues, enabling proactive management and maintenance.

Firmware and Software Updates:

- Simplifies the process of updating firmware and software across multiple devices.
- Ensures infrastructure is always running the latest, most secure versions.

Integration and APIs:

- Integrates with popular IT management and orchestration tools, allowing for seamless workflows and processes.
- Provides REST APIs for custom integration and automation.

Resource Management:

- Helps in capacity planning and ensuring efficient use of infrastructure resources.

11. Appendix:

11.1 Hardware components

Table 6: ThinkEdge SE360 V2 Hardware Component

Component	Specification
Form factor	ThinkEdge, 2U high
Processor	One Intel Xeon D-2700 Series processor (formerly codenamed "Icelake D"). Supports processors up to 16 cores, core speeds of up to 2.1 GHz, and TDP ratings of up to 100W. The processor is soldered onto the system's board.
Memory maximum	Up to 256GB with 4x 64GB RDIMMs
Security features	TPM 2.0 (WW) / TCM Nation Z (China only).
Mounting options	Standalone, desktop mount, DIN Rail or Wall mount, ceiling mount and Standard Rack (2U2N).
Operating systems supported	Alma Linux, Canonical Ubuntu, Ubuntu enterprise, Microsoft Windows Server, Red Hat Enterprise Linux, Rocky Linux, SUSE Linux Enterprise Server, VMware ESXi. See the Operating system support section for specifics.
Cooling	Two non hot-swap 80mm FANs, N+1 redundant in most configurations.
Temperature	Handles temperature changes from 0-55°C as well as high dust and vibration environments. Also supports extreme extended temperatures -20~65°C with certain configurations and with marine certification.



Figure 5: ThinkEdge SE360 V2

Table 7: ThinkEdge SE450 V2 Hardware Components

Component	Specification
Form factor	2U rack server, 300mm or 360mm depth
Processor	1x 3rd Gen Intel Xeon Scalable processor (formerly codenamed "Ice Lake"). Supports Silver, Gold and Platinum processors up to 36 cores, core speeds of up to 3.0 GHz, and TDP ratings of up to 205W.
Memory maximum	With RDIMMs only: Up to 1TB with 8x 128GB 3DS RDIMMs With PMem: Up to 1.5TB with 4x 128GB 3DS RDIMMs + 4x 256GB PMem modules
Security features	ThinkShield Key Vault Portal web site for security management. Trusted Platform Module, supporting TPM 2.0. Lockable top cover, optional front locking security bezel with duplicate status LEDs, intrusion detection, self-encrypting drive (SED) support, power-on password, administrator's password.
Mounting options	Horizontal or vertical orientation. Floor mount (vertical), wall mount, or 2U short-depth rack mount. Available locking security bezel with dust filter.
Operating systems supported	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, VMware ESXi. See the Operating system support section for specifics.

Cooling	6x simple-swap 60 mm dual-rotor fans (all standard), N+1 redundant in most configurations.
Temperature	Handles continuous operating temperatures from 5°C and 45°C; and some configurations are designed to meet NEBS Level-3 and ETSI requirements for 96 hours operating excursions from -5°C and 55°C as well as tolerance to locations with high-dust and vibration




Figure 6: ThinkEdge SE450

Table 8: ThinkEdge SE360 V2 Hardware Components

Component	Specification
Form factor	2U rack server, short depth (438mm depth, from EIA front rack flange)
Processor	1x AMD EPYC 8004 processor (formerly codenamed "Siena"). Supports up to 64 cores, core speeds of up to 2.65 GHz, and TDP ratings of up to 225W.
Memory maximum	Up to 576 GB with 6x 96GB RDIMMs

GPU and accelerators	Supports up to 6x single-wide GPUs and accelerators Supports up to 2x double-wide GPUs
Security features	Centralized security management with ThinkShield Key Vault Portal web site or manage security directly in XCC2. Platform Firmware Resiliency (PFR) Root of Trust (RoT), Trusted Platform Module, supporting TPM 2.0. Lockable top cover, intrusion detection, self-encrypting drive (SED) support, power-on password, administrator's password. Optional front locking security bezel with duplicate status LEDs
Mounting options	2-post rack, 4-post short-depth 600mm rack, 4-post 1000mm rack. Available locking security bezel with dust filter. Available security EIA brackets.
Operating systems supported	Microsoft Windows Server, Red Hat Enterprise Linux, SUSE Linux Enterprise Server, VMware ESXi. See the Operating system support section for specifics.
Cooling	5x simple-swap 60 mm dual-rotor fans (all standard), N+1 rotor redundant.
Temperature	Handles continuous operating temperatures from 5°C and 55°C as has tolerance to locations with high-dust and vibration.



Figure 7: ThinkEdge SE455 V3

11.2 Lenovo Bill of materials

This appendix contains the bill of materials (BOMs) for different configurations of hardware for the Public Sector Kognition Safety and Security Edge AI solution deployment. The BOM (Bill of Materials) lists in this appendix are not meant to be exhaustive and must always be double-checked with the configuration tools. Any discussion of pricing, support, and maintenance options is outside the scope of this document.

Table 9: Numbers of servers needed for different camera counts

Extra Small Size - Product Description	QTY
SE360 XS 8C 64GB 1TB QC: ThinkEdge SE360 V2 with Controlled GPU	1
ThinkEdge SE360 V2 Chassis	1
ThinkEdge SE360 V2 4x 1Gb, 2x 2.5Gb (TSN) I/O Module	1
System Operational Temperature 0C to 55C / 32F to 131F	1
Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
ThinkEdge SE350 V2/SE360 V2 Planar with Intel Xeon D-2733NT 8C 80W 2.1 GHz	1
ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	2
Select Storage devices - no configured RAID required	1
On Board SATA Software RAID Mode	1
Intel VROC (VMD NVMe RAID) Premium	1
ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	2
ThinkEdge SE360 V2 M.2 Cabled Adapter Module	1
Select Storage Devices - Configured M.2/7mm RAID	1
M.2 NVMe Array 1 RAID 1	1
ThinkSystem M.2 N-30m2 480GB Read Intensive NVMe PCIe 3.0 x4 NHS SSD (Industrial)	2
ThinkSystem Qualcomm Cloud AI 100	1
ThinkEdge SE360 V2 Riser Assembly (M.2 Riser + PCIe Riser)	1

Internal AC Power Supply (230V/115V)	1
ThinkEdge SE360 V2 AC Power Input Board	1
ThinkEdge SE360 V2 500W 230V/115V Non-Hot Swap Power Supply	1
2.8m, 10A/120V, C13 to NEMA 5-15P (US) Line Cord	1
ThinkEdge SE360 V2 Fan Assembly (Front to Rear)	1
Small Size - Product Description	
SE360 S 12C 128GB 1TB QC: ThinkEdge SE360 V2 with Controlled GPU	1
ThinkEdge SE360 V2 Chassis	1
ThinkEdge SE360 V2 4x 1Gb, 2x 2.5Gb (TSN) I/O Module	1
System Operational Temperature 0C to 55C / 32F to 131F	1
Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
ThinkEdge SE360 V2 Planar with Intel Xeon D-2752TER 12C 77W 1.8 GHz	1
ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM Extended Temperature	4
Select Storage devices - no configured RAID required	1
On Board SATA Software RAID Mode	1
Intel VROC (VMD NVMe RAID) Premium	1
ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	2
ThinkEdge SE360 V2 M.2 Cabled Adapter Module	1
Select Storage Devices - Configured M.2/7mm RAID	1

M.2 NVMe Array 1 RAID 1	1
ThinkSystem M.2 N-30m2 480GB Read Intensive NVMe PCIe 3.0 x4 NHS SSD (Industrial)	2
ThinkSystem Qualcomm Cloud AI 100	1
ThinkEdge SE360 V2 Riser Assembly (M.2 Riser + PCIe Riser)	1
Internal AC Power Supply (230V/115V)	1
ThinkEdge SE360 V2 AC Power Input Board	1
ThinkEdge SE360 V2 500W 230V/115V Non-Hot Swap Power Supply	1
2.8m, 10A/120V, C13 to NEMA 5-15P (US) Line Cord	1
ThinkEdge SE360 V2 Fan Assembly (Front to Rear)	1
Small plus Size - Product Description	
SE360 S+ 16C 128GB 1TB QC: ThinkEdge SE360 V2 with Controlled GPU	1
ThinkEdge SE360 V2 Chassis	1
ThinkEdge SE360 V2 4x 1Gb, 2x 2.5Gb (TSN) I/O Module	1
System Operational Temperature 0C to 55C / 32F to 131F	1
Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
ThinkEdge SE350 V2/SE360 V2 Planar with Intel Xeon D-2775TE 16C 100W 2.0 GHz	1
ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	4
Select Storage devices - no configured RAID required	1
On Board SATA Software RAID Mode	1
Intel VROC (VMD NVMe RAID) Premium	1

ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	2
ThinkEdge SE360 V2 M.2 Cabled Adapter Module	1
Select Storage Devices - Configured M.2/7mm RAID	1
M.2 NVMe Array 1 RAID 1	1
ThinkSystem M.2 N-30m2 480GB Read Intensive NVMe PCIe 3.0 x4 NHS SSD (Industrial)	2
ThinkSystem Qualcomm Cloud AI 100	1
ThinkEdge SE360 V2 Riser Assembly (M.2 Riser + PCIe Riser)	1
Internal AC Power Supply (230V/115V)	1
ThinkEdge SE360 V2 AC Power Input Board	1
ThinkEdge SE360 V2 500W 230V/115V Non-Hot Swap Power Supply	1
2.8m, 10A/120V, C13 to NEMA 5-15P (US) Line Cord	1
ThinkEdge SE360 V2 Fan Assembly (Front to Rear)	1
Medium Size – Product Description	
SE450 M 32C 128GB 4TB QC: ThinkEdge SE450 - 3yr Warranty with Controlled GPU	1
ThinkEdge SE450 Edge Server Chassis, 300mm	1
Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
Intel Xeon Platinum 8352Y 32C 205W 2.2GHz Processor	1
ThinkSystem SR650 V2 Standard Heatsink	1
ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	4

Select Storage devices - no configured RAID required	1
Intel VROC (VMD NVMe RAID) Premium	1
ThinkSystem 2.5" U.2 7mm PM9A3 1.92TB Read Intensive NVMe PCIe 4.0 x4 Trayless SSD	3
ThinkEdge SE450 Internal HDD Cage	1
Select Storage Devices - Configured M.2/7mm RAID	1
ThinkSystem M.2 SATA/NVMe 2-Bay Adapter	1
On Board SATA Software RAID Mode for M.2	1
M.2 SATA Array 1 RAID 1	1
ThinkSystem M.2 ER3 480GB Read Intensive SATA 6Gb NHS SSD	2
ThinkEdge Intel I350 1GbE RJ45 4-port Internal Lock OCP Ethernet Adapter	1
ThinkSystem Qualcomm Cloud AI 100	2
ThinkEdge SE450 Riser 2	1
ThinkEdge SE450 Riser 1	1
ThinkEdge 1100W 230V/115V Platinum Hot-Swap Power Supply	2
Large Size Product Description	
SE455 L 64C 256GB 4TB QC: ThinkEdge SE455 V3 - 3Yr Warranty with Controlled GPU	1
ThinkEdge SE455 V3 Chassis	1
Operating mode selection for: "Maximum Efficiency Mode"	1
Acoustic Mode Disabled	1
System Operational Temperature 5C to 25C / 41F to 77F	1

ThinkEdge SE455 V3 AMD EPYC 8534P 64C 200W 2.3GHz Processor	1
ThinkEdge SE455 V3 2U Heatsink	1
ThinkSystem 64GB TruDDR5 4800MHz (2Rx4) 10x4 RDIMM-A	4
Select Storage devices - no configured RAID required	1
ThinkSystem RAID 5350-8i PCIe 12Gb Adapter	1
ThinkSystem 2.5" S4520 1.92TB Read Intensive SATA 6Gb HS SSD	3
ThinkEdge SE455 V3 4x2.5" SAS/SATA Backplane	1
ThinkEdge SE455 V3 M.2 SATA/x4 NVMe Adapter with Carrier	1
Select Storage Devices - Configured M.2/7mm RAID	1
ThinkSystem RAID 5350-8i for M.2 SATA Boot Enablement	1
M.2 SATA Array 1 RAID 1	1
ThinkSystem M.2 5400 PRO 480GB Read Intensive SATA 6Gb NHS SSD	2
ThinkSystem Broadcom 5719 1GbE RJ45 4-port OCP Ethernet Adapter	1
ThinkSystem Qualcomm Cloud AI 100	2
ThinkEdge SE455 V3 Riser1	1
ThinkEdge SE455 V3 Riser2	1
ThinkEdge 1100W 230V/115V Platinum Hot-Swap Power Supply	2

11.3 Abbreviations

Table 10: Abbreviations

Acronym	Meaning
---------	---------

ISV	Independent Software Vendor
OS	Operating System
GUI	Graphic User Interface
AI	Artificial Intelligence
LVD	Lenovo Validated Design
GPU	Graphic Processing Unit
PTZ	Pan Till Zoom
HTML	Hypertext Markup Language
RAM	Random Access Memory
VM	Virtualized Machine
RTSP	Real-Time Streaming Protocol
IP Camera	Internet Protocol Camera
SSD	Solid State Disk

Resources

Table 11: Resources

Resources	Links
Kognition	Kognition – Connect – Kognition Inc
SE360 V2	Lenovo ThinkEdge SE360 V2 Server Product Guide > Lenovo Press
SE450	Lenovo ThinkEdge SE450 Edge Server Product Guide > Lenovo Press
SE455 V3	Lenovo ThinkEdge SE455 V3 Server Product Guide > Lenovo Press
Capgemini	2024 Key Trends in Public Security
Expert Market Research	Public Kognition Safety and Security Market Size, Analysis 2024-2032
Genetec	Physical security: Trends for 2024
La Jolla Country Day School Case Study	https://www.lenovo.com/us/en/case-studies-customer-success-stories/ljcds

Document history

Version 1.0	July 2024	First version includes Lenovo 360 V2, SE450, and SE455 V3 with Kognition AI Software
-------------	-----------	--

Trademarks and special notices

© Copyright Lenovo 2024.

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®
- ThinkShield®
- ThinkSystem®
- TruDDR4
- 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.

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

Microsoft®, Windows Server®, and Windows® are trademarks 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 depending on the 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.

The 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.