Lenovo

Lenovo Validated Design for Al-Powered Industrial Workplace Safety with Avathon

Last update: 06 May 2025

Version 1.0

Automatically detect PPE non-compliance and workplace hazards in real time Validated on Lenovo ThinkEdge for scalable, secure, and highperformance deployments in industrial environments

Run AI at the edge with cloudbased training for continuous model improvement Easily integrate with existing video systems for Rapid Deployment & Interoperability

Vanita Meyer Sinan Atan



Table of Contents

Introduction	1
Business Problem and Business Value	2
Business Problem	2
Business Value	2
Technology Overview	4
Functional Requirements	4
Non-functional requirements	5
Architectural overview	7
Component model	10
VAIA Service Modules	11
ThinkEdge Platform	13
ThinkEdge SE360 V2 ThinkEdge SE455 V3	
Operational Model	
Avathon VAIA Operational Model	
Lenovo ThinkEdge Systems Operational Model	19
Deployment considerations	21
Systems management	21
Server / Compute Nodes	21
Networking	22
Storage integration	22
Sizing Considerations	22
Solution Validation Results	23
Validation Methodology	23

Summary of Validation Results	24
Solution Summary	25
Appendix: Lenovo Bill of Materials	26
Appendix: Abbreviations	29
Resources	32
Document history	33
Trademarks and special notices	34

Introduction

This Lenovo Validated Design (LVD) delivers a turnkey, AI-powered solution for improving workplace safety in industrial environments. Developed in partnership with Avathon, a leader in visual AI technologies, this solution empowers organizations to monitor, detect, and prevent safety hazards in real time. Leveraging Lenovo's AI-optimized edge infrastructure and Avathon's Visual AI Advisor (VAIA) platform, the design offers pre-tested hardware and software integration to ensure faster time to value, simplified deployment, and regulatory compliance.

The intended audience includes IT decision-makers, safety and compliance officers, operations managers, and AI solution architects looking to enhance safety, reduce operational risks, and accelerate digital transformation in manufacturing, energy, logistics, and other heavy industries.

Business Problem and Business Value

Business Problem

Across industrial sectors, workplace safety continues to be a top priority—yet remains a persistent challenge. Organizations struggle with real-time monitoring of safety violations and proactively identifying risks that can lead to costly incidents, injuries, and regulatory penalties. Traditional safety methods, such as manual inspections or reactive incident reporting, are labor-intensive, error-prone, and insufficient in dynamic, highrisk environments.

At the same time, industries are under growing pressure to meet increasingly stringent health and safety regulations (e.g., Occupational Safety and Health Administration (OSHA), Health, Safety and Environment (HSE), Application Programming Interface (API)) while also minimizing downtime and insurance costs. The lack of integrated, scalable, and intelligent safety systems prevents companies from effectively addressing these challenges.

This LVD combines Lenovo's rugged, high-performance ThinkEdge infrastructure with Avathon's advanced Visual AI Advisor (VAIA) platform to deliver a scalable edge AI solution tailored for workplace safety. The solution transforms existing video infrastructure, such as closed-circuit television (CCTV) systems, drones, and mobile cameras—into real-time safety monitoring tools. It detects Personal Protective Equipment (PPE) compliance (e.g., helmets, vests), identifies hazardous behaviors (e.g., person down, unsafe vehicle proximity), and alerts operators instantly to potential risks.

Business Value

The Lenovo Validated Design (LVD) for AI-Powered Workplace Safety delivers tangible business value by helping industrial organizations protect their workforce, minimize operational risks, and maintain compliance with evolving safety regulations. At the core of this solution is Avathon's Visual AI Advisor (VAIA), an advanced visual intelligence platform designed to deliver real-time hazard detection and automated safety enforcement through AI and computer vision.

Avathon's solution leverages state-of-the-art AI models to detect high-risk scenarios and unsafe behaviors, including PPE non-compliance, worker falls (person down), vehicle near-miss incidents, pathway obstructions, fire and smoke, spill detection, and unauthorized access. These events are identified in real time, enabling immediate alerts and corrective action. By shifting safety monitoring from manual and reactive to automated and proactive, the solution helps reduce workplace injuries, downtime, and costly safety violations.

A key differentiator is VAIA's seamless integration with existing video infrastructure, including CCTV systems, drones, and mobile cameras—allowing rapid deployment without major hardware overhauls. The platform runs efficiently on Lenovo ThinkEdge servers, enabling real-time AI processing directly at the edge. This local processing architecture eliminates the need for cloud dependency, ensuring ultra-low latency, high system availability, and enhanced data privacy—critical for regulated industrial environments.

The solution's modular, container-based architecture, deployed within a hybrid cloud environment, allows centralized management, remote operations, and scalable expansion across facilities. Built-in capabilities for AI model retraining ensure continuous learning and improving detection accuracy as workplace conditions evolve. Through its intuitive dashboards and mobile accessibility, safety teams gain full visibility and control, improving both responsiveness and long-term strategic safety planning.

Coupled with Lenovo's rugged, enterprise-grade ThinkEdge systems, this solution ensures high performance under demanding conditions. Lenovo provides pre-tested and optimized hardware configurations to match different deployment scales, small, medium, or large, ensuring right-sized investment and resource efficiency. With validated benchmarks, interoperability testing, and sizing guidelines included, this LVD simplifies adoption and reduces time to value.

Ultimately, this solution enables organizations to foster a safety-first culture powered by intelligent automation, enhancing compliance, reducing liability, and improving workforce productivity through smarter, faster, and more resilient safety operations.

Technology Overview

Functional Requirements

Avathon Visual Al Advisor (VAIA) is an advanced computer vision platform designed to transform traditional video surveillance into actionable intelligence for enterprise environments. By leveraging Al-powered analytics, VAIA helps organizations enhance safety, security, and operational efficiency through real-time monitoring and automated insights. Designed for seamless integration, customization, and deployment, Avathon's platform is revolutionizing industrial AI.

The following is a list of functional requirements delivered as part of the VAIA solution:

- **Problem Definition & Business Case Development:** A structured methodology to identify improvement opportunities, connect insights to operational changes, and build a compelling business case for Computer Vision AI adoption and investment.
- **Single Stack:** Use cases enablement via a single stack software rapidly accelerating time-to-deploy for new scenarios in case of need with minimal investment.
- **AI-Powered Computer Vision:** Utilizes advanced imaging technology to capture operational behaviors and analyze them against predefined criteria to deliver actionable insights to HSE and operational quality assurance teams—supporting smarter decisions and improved efficiency.
- Seamless Enterprise Integration: Facilitates closed-loop automation by integrating with enterprise systems via API (software) or MQTT (hardware), supporting automation, streamlining workflows, and driving operational efficiency.
- **Scalable Architecture:** A flexible, AI-enabled solution designed to integrate effortlessly across enterprise environments, ensuring compatibility between hardware and software components.
- **Business-Focused Strategy:** Emphasizes a strategic approach to initial deployment and scaling, maximizing return on investment in health, safety, security, and productivity related operations.
- Real-Time Video Analytics
 - **Continuous Monitoring**: VAIA processes live video feeds 24/7, enabling immediate detection of safety violations, security breaches, and operational inefficiencies.
 - **Pre-Trained Use Cases**: Supports over 125 pre-built scenarios, including PPE compliance, unauthorized access detection, and equipment productivity.
- Alerting and Notification System
 - Multi-Channel Alerts: Delivers notifications through web dashboards, mobile apps, email, SMS, WhatsApp, and on-site alarms and machine integrations, ensuring prompt response to incidents.
 - Integration with Communication Systems: Allows third-party systems to consume data and insights via APIs and MQ Telemetry Transport (MQTT), facilitating seamless communication across platforms.

Non-functional requirements

In addition to its core features and functions, the solution also delivers non-functional capabilities that ensure reliable and efficient execution of use cases in edge environments.

- Integration with Existing Infrastructure
 - Hardware Compatibility: Works seamlessly with current CCTV systems, Pan-Tilt-Zoom (PTZ) cameras, fixed-view bullet cameras, multi-sensor cameras, drones, and mobile devices, eliminating the need for new hardware investments via Real Time Streaming Protocol (RTSP) streams¹.
 - **IT Systems Integration**: Standardized APIs for integration with enterprise systems (Enterprise Resource Planning (ERP), Manufacturing Execution System (MES), etc.)
 - **Deployment Flexibility**: Can be implemented on-premises or via cloud services, utilizing cost-effective Graphics Processing Unit (GPU) nodes for scalability [6].
- **Compact Form Factor:** Optimized for edge environments, the solution operates without the need for a dedicated data center.
- User-Friendly Interface
 - No-Code/Low-Code Platform: Features a drag-and-drop web interface and mobile app, allowing users to configure and manage the system without programming skills.
 - Customizable Dashboards: Provides role-based alerts and analytics tailored to specific operational needs.
- Performance & Throughput
 - o Real-time processing with minimal latency
 - o Handles large volumes of video data from multiple sources
- Scalability and Adaptability
 - **Rapid Deployment**: Capable of scaling to thousands of cameras within days, consolidating all video feeds into a single interface.
 - Continuous Learning: Incorporates feedback to improve performance over time by reducing false positives and negatives.
- Security and Compliance
 - Data Security: Ensures secure data processing, whether on-site or in the cloud, maintaining compliance with industry standards. Encrypted communication between users and the UI is needed. When hosted in ISV's private cloud, disk encryption for encryption at REpresentational State Transfer (REST) is required.

¹Fisheye and high-distortion view cameras are not supported

• Secure Integration: Strong security protocols are in place to protect customer data and infrastructure from threats and vulnerabilities. Encryption support in transit using application gateways is required to protect data between different horizontal or vertical layers. All platform deployed platform components shall be scanned at build time using WhiteSource for vulnerabilities and license dependencies.

- **Compliance Monitoring**: Monitors adherence to safety protocols and regulations, aiding in achieving Health, Safety, and Environmental (HSE) goals
- **Fault Tolerance:** Designed to maintain functionality even when individual components fail, ensuring consistent uptime and reliability.
- Easy Installation: The deployment process is simplified to reduce setup time and complexity.
- **Streamlined Management:** Built for straightforward monitoring and maintenance, the solution supports efficient daily operations with modular architecture. It supports auto-updates, monitoring and built-in diagnostic tools via automation.

Architectural overview

The solution is deployed on Lenovo ThinkEdge servers, which handle AI inferencing at the edge, while the training engine and application layer operate in a public cloud environment. A secure connection is established between the private (edge) and public (cloud) environments using one of the following approaches:

- Option 1, A unique identity key is provided for each device which is used by the device to communicate with the Internet of Things (IoT) Hub.
- Option 2, A X509 certificate and a private key is combined to authenticate the device to the IoT Hub.

The security token method provides authentication for each call made by the device to IoT Hub by associating the symmetric key to each call. X.509-based authentication allows authentication of an IoT device at the physical layer as part of the Transport Layer Security (TLS) connection establishment. The security-token-based method can be used without the X.509 authentication which is a less secure pattern. The choice between the two methods is primarily dictated by how secure the device authentication needs to be, and availability of secure storage on the device (to store the private key securely).

The keys should not be sent over the network and hence security tokens are used to authenticate devices and services. These tokens have a limited validity and need to be renewed. While using the Azure SDK the tokens are regenerated automatically. In the data plane, information is collected from far-edge devices—including sensors, cameras, and Network Video Recorder/Digital Video Recorder (NVR/DVR) systems, and ingested either via streaming protocols such as RTSP and User Datagram Protocol (UDP). This data is then processed by GPU-powered ThinkEdge devices, which generate telemetry data, alarms, and events formatted in JavaScript Object Notation (JSON) over Hyper-Text Transfer Protocol/Hyper-Text Transfer Protocol Secure (HTTP/HTTPS). These outputs are forwarded to the public cloud for further processing at the application layer and for distribution to subscribers. In the case of hardware and device automation, an MQ Telemetry Transport (MQTT) broker is used to integrate with hardware and physical systems.

At deployment, video clips are sent to the training engine in the cloud where Avathon will incorporate new data to tune models. This engine continuously improves and optimizes the AI models for the specific deployment environment. The updated models are then distributed back to the edge for inferencing via cloud-to-edge configuration sync.

The management layer uses HTTPS communication between the public cloud and edge environments for secure container registry access and orchestration of application deployment, ensuring efficient management and updates of edge workloads. Internet connection between the edge and cloud is secured using the TLS standard. VAIA supports TLS 1.2, TLS 1.1 and TLS 1.0, in this order. Support for TLS 1.0 is provided for backward compatibility only. It is recommended to use TLS 1.2 since it provides the most security.

Figure-1 provides an overview of deployment Architecture of the solution.



Figure1– Solution Architecture Overview

Table 1 outlines the technologies validated in each of the architecture layers as part of this solution:

Layer	Technology	Version	Role
Far Edge (Cameras)			30fps, 1920x1080, 2048kbps
Edge Infrastructure	ThinkEdge	SE360 V2 SE455 V3	SE360 V2 for small & medium; SE455 V3 for large models
Operating System	Ubuntu 22.04	22.04	Long-Term Support (LTS) version
CaaS	Docker	28.0.4	Container management (K8, etc.) as
	Helm	3.17.3	option for multi-modal deployments
	Kubernetes	1.32.3	
Drivers	Nvidia	V550	GPU Drivers
Inferencing & Runtime	Avathon VAIA	3.5.0	Real-time AI model execution
Model Training	Avathon VAIA	3.5.0	AI model retraining for improved
			accuracy and adaptability
Data Lake	Avathon VAIA	3.5.0	Centralized storage for analytics and model training

Table 1 – LVD Technology Mapping

Insights/dashboard Avathon VAIA	3.5.0	Data driven decision making
---------------------------------	-------	-----------------------------

Component model

Avathon's AI Platform is structured into four interrelated layers, each playing a critical role in enabling intelligent, responsive, and scalable operations. These layers work together internally and interface with the external environment to deliver a comprehensive AI-driven solution:

Data Foundation

The Data Foundation layer interfaces with far-edge devices and various data sources. It is responsible for integrating, managing, and preparing the data required to train, retrain, and drive inferencing for specific use cases.

AI Core Services:

Al Core Services Layer are foundational Al model inferencing services for processing real-time data received from the data foundation layer to produce Al insights with the help of GenAl, Natural Language Processing (NLP), Deep Learning, and Computer Vision services.

Application Programming Interfaces

Built on REST-based, open standards, these interfaces connect the AI platform with external systems. They facilitate the delivery of actionable insights by enabling event-driven notifications, real-time monitoring, predictive alerts, and prescriptive actions.

Autonomous Operations

This layer drives automation by triggering workflows based on the outputs from the AI Core Services. It supports activities such as raising alerts, initiating provisioning tasks, and notifying administrators, enabling efficient and autonomous system management.

Figure 2 provides a brief overview of these layers:



VAIA Service Modules

Scope of this LVD is limited to following use cases:

- Intrusion
- Fire Detection
- Spill Detection
- Forklift/Vehicle Near Miss
- Pathway Compliance
- Personal Protective Equipment (PPE)
- Person down

These use cases run simultaneously during validation executing computer vision AI model to generate results on each individually. Further details about the testing can be found in 'Solution Validation Results' Section of the document.

Table below outlines the complete list of application services can be activated in the software in diverse verticals:

#	Category	Use Case
1		Crowd Monitoring
2		Detect Doors Left Open
3		Person Down
4	Surveillance & Monitoring	Blocked Pathways
5		Intrusion-Person, Vehicle
6		Loitering
7	- Vehicle Management	Vehicle Count & Type
8		Detect Repeat Visits
9		Queued Vehicles
10		Parking Duration
11		Vehicle Speeding
12		Unauthorized Parking

13		Forklift Near Miss
14	Workplace Hazard Detection	Machine Guarding
15		Person Down
16		Pathway Compliance
17		Spillage
18		Working under Suspended Load
19		Detect Fire & Smoke
20	Emergency Response	Person Down
21	PPE Compliance	Detect PPE – Hair Net, Gloves, Helmets, Lab Coats
22	Optimize Storage Operations	Inventory Management
23		Checkout Queue Management
24		Improve Employee Productivity
25		Optimize Restroom Cleaning
26		Detect Drops Off Outside Delivery Times
27		Foot Traffic Analysis
28		Identify Demographic Trends
29	Customer Analytics	Identify High Activity Areas
30		Time Spend at the Facility
31		Visitor to Customer Conversations, Loss of Business
32		Customer Journey
33		Dock Door Utilization
34	Yard Management	Vehicle Tracking
35		Vehicle Speeding
36	Weapons Detection	Detect Person Holding Rifles & Handgun

37	- Material Handling	Derailing on Conveyor
38		Colliding Packages
39		Package Count
40		Falling Package Count

ThinkEdge Platform

Lenovo's ThinkEdge portfolio provides data center-like computing at the edge ensuring the flexible way to power innovation for AI powered services. ThinkEdge devices offer tailored solutions for improved performance, security and scalability. Edge servers are rugged and secure with physical tamper-proofing, data encryption and the ability to withstand conditions of all kinds. In a nutshell ThinkEdge achieves:

- **A Hybrid Multi-Cloud with Flexibility** The Lenovo Edge ecosystem is open and flexible. It integrates with all key cloud providers to give end-to-end, ongoing management.
- Extend Cloud with a Resilient Edge an integrated, scalable solution that's simple to deploy, costeffective to operate, and built to maximize performance for edge workloads.
- **Improve Data Management and Access** Improve data management and accessibility with a wide range of solutions and expert guidance tailored for modern hybrid infrastructure.

ThinkEdge SE360 V2

Lenovo's ThinkEdge SE360 V2 is the most rugged AI server being one of the best options for inferencing at the far edge. As part of this LVD SE360 V2 is scoped with Phase 1 for inferencing purposes. Below is the key characteristics of SE360 V2:

- Intelligent System Cooling Optimized for quiet operations in occupied spaces & dust filtering for ultimate reliability being twice as quiet compared to all other competitive Edge servers on the market
- Broad Networking, Input/Output (I/O) and GPU Expansion for integrating control systems and deploying AI being the smallest edge AI server on the market, rick networking options with Wireless Local Area Network (WLAN) support and up to 100gbE ethernet connection
- Energy Efficient Configurable Energy Optimized Mode for workloads performance to reduce energy consumption
- Enhanced Security Lockable bezel with smart filtering technology supported with disk encryption, ThinkShield Activation, system lockdown, movement detection and tamper protection
- Compact Form Factor 78% smaller than competitive products for easy deployment
- **Data Protection** If the system is moved or tampered, beyond the configuration threshold, the data is automatically lock-up
- Unique Mounting Options Short-depth for ultimate space savings and secure cable routing
- Shock & Vibration Resistant Designed to perform with high level of shock & vibration up to 40G

Form Factor	2U Height, Width: 212mm, Depth: 317.5mm
Processor	Intel® Xeon® D processor, up to 20 cores, up to 125W Thermal Design Power (TDP)
Drive Bays (2 Internal)	Up to 2 x SATA/NVMe 2.5-inch 7mm drives HS; Up to 8 x M.2 using PCIe slots 2 x M.2 boot drives (Redundant Array of Independent Disks (RAID) 1) + 2 x M.2 Data drives internally
Memory	4 x DDR4 memory slots; Maximum 256GB
Expansion Slots	1 x PCle G4 x16 HHHL + 1 x PCle G4 x 16 FHHL
GPUs	Nvidia A2 and L4, Qualcomm Cloud AI 100, Intel Flex 140
Network Interface Front panel & I/O	 4 x 10GbE/25GbE SFP+/SFP28 2 x 2.5GbE Time Sensitive Network (TSN) 1GbE RJ45 management port (switchable WLAN-Gen6) or 4 x GbE 2 x 2.5GbE Time Sensitive Network (TSN) 1GbE RJ45 management port (switchable WLAN-Gen6) Power button LED, ID button LED, BT enabling button LED (BT enable request, BT active), Status (System error) LED, Activation LED, BMC WLAN Active LED, x86 WLAN Active LED, 2 x USB 3.0 (Type-A) + 1 x BMC USB 2.0 (Type-A), USB-C (USB 2.0 + Display Det (widee) (USB 2.0 auto avitab) PMC actival
HBA/RAID Support	RAID 0, 1, 5, 10 (with Intel VROC (Virtual RAID on CPU) RAID support)
Power	Single power supply AC 550W DC power supplies 12V-48V
Security	ThinkShield Activation, Tamper protection, Motion detection, Encrypted SSD, System Lockdown, GPS tracking, TPM 2.0
Easy to Deploy	ThinkShield activation, Lenovo Open Cloud Automation (LOC-A)

Table 1 - ThinkEdge SE360 V2 Characteristics

Operating temperature	-20°C to 65°C
Systems Management	Lenovo XClarity Controller (AST2600)
OS Support	Microsoft, Red Hat, Ubuntu, SUSE, VMware
Limited Warranty	3-year customer replaceable unit and onsite service, next business day 9x5; optional service upgrades

ThinkEdge SE455 V3

The Lenovo ThinkEdge SE455 V3 is a high-performance edge server built for demanding AI workloads in industrial and remote environments. It offers exceptional compute density, energy efficiency, and rugged reliability, making it ideal for GPU-intensive training, data aggregation, and real-time analytics outside the traditional data center.

- Short depth chassis for use outside the data center 2U short-depth form factor with front access I/O ensuring dust filtering, extended temperature operation (5 to 55 C), shock & vibration resistance
- Very flexible I/O and storage configuration to support most demanding edge applications Optional internal drive bays, optional second PCIe riser, high speed Gen5 x16 Open Compute Project (OCP) achieving 2x more storage bays delivering over 490 TB of total storage, 20% more PCIe slots than competitive edge servers
- Quiet Operation Acoustic modes, enhanced heat dissipation via noise controlling with Unified Extensible Firmware Interface (UEFI) settings, larger Central Processing Unit (CPU) heatsink
- Energy Efficient Advanced Micro Devices (AMD) Edge optimized processers fine-tuned with UEFI operating modes resulting in up to 64 cores in socket achieving 32% better energy efficiency.
- **GPU & Memory Rich** Up to 6x single-width GPUs or 2x double-width GPUs supported with up to 768GB of system memory
- Increased security to protect data Encrypted disk and ThinkShield activation achieving tamper protection and system lockdown in case of security incidents

Form Factor	2U rack server 440mm depth
Processor	1x AMD EPYC 8004 Series Processor, up to 64 cores, up to 225W TDP
Drive Bays	Up to 4x 2.5-inch Front Hot Swap 15 mm drives Up to 4x 2.5-inch Internal 15 mm drives (optional) 2x M.2 boot drives: non-RAID or Hardware M.2 RAID
	Common options with ThinkSystem

Table 2 - ThinkEdge SE455 V3 Characteristics

Memory	6x DDR5 memory slots, up to 576GB using 96GB RDIMM (128GB DIMM in plan) Common options with ThinkSystem
Expansion Slots	2x PCle Gen5 x16 (Riser2 optional) 4x PCle Gen4 x8 (Riser2 optional) 1x OCP 3.0 Gen5 x16
GPU/Accelerator	Standard Support: Up to 6x NVIDIA A2, L4, L40, AMD V70, Qualcomm AI 100 Core Supported: AMD Xilinx T2, Intel ACC100, others on request Future Support: NVIDIA L40S, others on request
Network Interface	1/10/25/100 Gb LOM adapter in OCP 3.0 slot Up to 6x 1/10/25/100/200 Gb PCIe network adapters
Ports, Buttons	Front: 1x Power Button (with green LED), 1x System Locator (with blue LED), 1x NMI button, 1x USB-C, 2x USB 3.0, 1x USB 2.0 for XCC2, 1x RJ45 for XCC2, 1x Diagnostic handset, COM port via PCI slot
LED	Security (green), Attention (yellow), XCC2 Ethernet (Link and Activity),
HBA/RAID Support	HW RAID with/without cache or SAS HBAs 4350-8i, 5350-8i, 440-8i, 540-8i, 940-8i with Supercap 2 and 4 port Qlogic QLE277x 32Gb Fibre Channel HBA
Power	Dual redundant power supplies AC (1100W Platinum/Titanium, 1800W Platinum) or Dual redundant power supplies -48V DC 1100W
Security	Security 2.0 with ThinkShield or XCC Management, Security Bezel, Tamper detection, Rack Security Bracket, Encrypted SSD, System Lockdown, Silicon Root of Trust, TPM 2.0, System Guard, AMD Infinity Guard, NIST SP800-193 compliance using hardware Root of Trust and Platform Firmware Resilience
Easy to Deploy	ThinkShield or XCC Managed lockdown mode and SED Lenovo Open Cloud Automation (LOC-A) and XClarity Administrator/Pro
Environmental	5°C to 55°C standard support ; NEBS 3 -5°C to 55°C (< 96 hours) 40/45dBA Acoustic Modes MERV2 Air Filter with clog detecting airflow sensor
Systems	Lenovo XClarity Controller (AST2600) DC-SCM 2.0

Management	
OS Support	Tier 1/1.5: Microsoft 2019/2022, Red Hat 8.x, 9.x, SLES 15, Ubuntu 20.x/22.x,VMware 7.x/8.x
	Tier 2: XenServer 8.2
	Tier 3: Alma Linux 8.x/9.x, Rocky Linux 8.x/9.x
	Future Support: Windows 11 IOT Enterprise Long-Term Servicing Channel (LTSC)
Limited Warranty	3-year customer replaceable unit and onsite service, next business day 9x5
	Optional service upgrades

Operational Model

Avathon VAIA Operational Model

Edge-Centric Inferencing

- The application is deployed on Lenovo ThinkEdge SE360 V2 and SE455 V3, which serve as the inferencing layer.
- These edge devices process video feeds in real time using AI models to detect anomalies, safety violations, and operational insights.
- Video input is streamed at 1080p resolution / 30 FPS from cameras, NVRs, or DVRs using RTSP or UDP protocols.

Cloud-Based Training and Application Layer

- Model training and re-training are performed in Avathon's private cloud, enabling continuous optimization based on environmental conditions and feedback loops.
- The application logic, dashboards, and APIs also reside in the cloud, providing users with centralized access to analytics, alerts, and configuration tools.

Secure Hybrid Connectivity

- An identity key or certificate and private key combination ensures secure communication between the edge and the cloud environments.
- Telemetry data, alerts, and configuration files are exchanged using HTTPS in JSON format.

Data Lifecycle and Automation

- Real-time data is used for model inferencing and provides immediate decision-making insights.
- **Autonomous operations** are triggered based on AI outputs (e.g., raising alerts, notifying staff, or provisioning resources) and may be integrated with physical systems via MQTT protocol.
- REST APIs support integration with enterprise systems for workflow automation, predictive maintenance, and compliance enforcement.

Management and Orchestration

- Edge workloads are orchestrated via cloud-based tools, using container registries and secure deployments over HTTPS. Solution configuration is managed through the web application.
- The platform supports multi-site scaling, centralized policy management, and system health monitoring.

This hybrid, closed-loop operational model ensures that VAIA delivers low-latency insights at the edge, while continuously evolving through cloud-based learning and orchestration, making it suitable for industries that demand both real-time intelligence and scalable AI deployment.

Lenovo ThinkEdge Systems Operational Model

Deployment orchestration and management of the stack can be done in container level integrating any validated solution in the environment. Lenovo Xclarity & LOC-A provides united dashboards to manage IT infrastructure and workloads with advanced automation capabilities both in hardware and Containers-as-a-Service (CaaS) layers. These tools also support open APIs, enabling integration with northbound systems to federate operations, streamline maintenance, and enhance end-to-end visibility.

Lenovo Xclarity

Lenovo XClarity software family is a set of foundation services for data center and edge infrastructures putting a premium IT execution capability and precision. Xclarity ensures:

- Standardized simple and automated infrastructure processes
- Freeing up time to deploy agile development methodologies and continual-delivery systems
- Ramping up delivery times of infrastructure and services.

XClarity software suite enables following modules:

- XClarity Controller It is the embedded management engine common in all ThinkEdge servers enabling fresh, uncluttered graphical user interface. XClarity Controller is Redfish-compliant webbased REST APIs for ease of inter-operability.
- XClarity Essentials Provides a suite of one-to-one management tools, including ONECLI, UpdateXpress, and Bootable Media Creator, to streamline system configuration, updates, and diagnostics.
- XClarity Administrator Management platform for ThinkEdge deployments scaling up to 1000 nodes. It exposes Graphical User Interface (GUI) capabilities for Hardware (HW) monitoring & Alerting, firmware updates, configuration and OS deployment. As well, it enables mobile app for anywhere management of infrastructure.
- XClarity Orchestrator Orchestrator enables management scalability for 10, 000 servers and hundreds of chassis for multiple domains easing administrative tasks under single glass-of-pane. It as well ensures analytics features as predictive failure information to support self-healing use cases lowering overall service costs.
- XClarity Energy Manager provides:

19

- o Monitoring capabilities to measure energy usage by device
- Optimization & planning to identify server utilization, power/thermal failures, energy capacity planning
- Energy efficiency automation for power-aware Virtual Machine (VM), container migration, power-aware job schedules and resilience in the presence of power outages.
- **XClarity Integrators** Integration into virtualization managers and upper-level management tools such as VMware, Microsoft, RedHat, Ansible, Chef, Puppet, Splunk, MS Azure Analytics, etc. are achieved via XClarity Integrators.

Lenovo Open Cloud Automation (LOC-A)

LOC-A helps enterprises to automate deployment, provisioning and management of platform services optimizing effort spend to administrate IT environments:

- Near-Zero Touch Provisioning Near-Zero Touch Provisioning (nZTP) enables scalable, latebinding deployment of edge servers, minimizing the need for field technician involvement. Instead of confirming early binding, LOC-A focuses on discovering what has arrived, with most tasks handled by the deployment admin via the LOC-A portal.
- OS Deployment LOC-A performs remote Operating System (OS) deployment on bare-metal nodes, automatically triggering and monitoring the process to ensure consistency and security across the fleet. This approach eliminates the need for Golden OS images in manufacturing or staging by provisioning the OS directly in the field. Additionally, LOC-A uses OS side-loading, transferring the OS image during provisioning, saving bandwidth by avoiding low-throughput WAN streaming.
- Plugin Mechanism for Partner Integration Lenovo Open Cloud Automation (LOC-A) features a plugin mechanism that allows Lenovo partners to easily create and integrate their own automated deployment flavours. This system leverages LOC-A's advanced features, enabling key functions for partner platforms, including:
 - o Bare-metal server provisioning and onboarding
 - o Orchestration of edge-node deployments
 - o Edge cluster and node instance creation (including OS deployment)
 - Access to a smart naming convention engine for hostname, Fully Qualified Domain Name (FQDN), etc. The Edge Platform uses this plugin mechanism to integrate seamlessly with LOC-A.
- Northbound API LOC-A also provides a secure, public northbound API for deeper integration with application orchestrators like Edge Platform or Operations Support System/Business Support System (OSS/BSS) platforms. This API enhances integration capabilities beyond the plugin mechanism, accelerating new feature enablement across platforms

Deployment considerations

A successful deployment of VAIA requires a strategic and structured approach to ensure smooth integration with existing IT systems. Key areas to address include system management, computing and server resources, networking, and data storage. This section highlights the essential deployment considerations:

- Performance tuning to support real-time processing
- Security protocols to safeguard data and infrastructure
- Scalability to support future expansion and evolving business requirements

Additionally, this section explains how the LVD underpins these capabilities, ensuring a dependable and resilient solution. It also provides a high-level overview of the IT environment requirements necessary for planning a successful VAIA deployment.

Systems management

A comprehensive and efficient systems management plan is vital to the continuous operation and maintenance of VAIA. Key focus areas for ensuring system reliability and security include:

- Monitoring and Performance Tracking Integrate VAIA status data with existing monitoring tools or use the built-in services to monitor system health, proactively identify issues, and gather detailed logs.
- Automated Updates and Security Patches Utilize automated tools to ensure the software environment stays updated with the latest features and security patches, minimizing downtime and vulnerabilities.
- **Remote Administration** Use Lenovo XClarity for remote management of ThinkEdge-based deployments, enabling improved operational control.

Server / Compute Nodes

VAIA *is* designed to support intensive AI workloads and is adaptable to centralized or edge computing environments depending on the use case. For peak performance, systems should be built on powerful computing nodes capable of handling real-time video analytics. Key considerations include:

- **Dedicated Al Hardware** Lenovo ThinkEdge servers equipped with high-performance CPUs, GPUs, and sufficient memory are ideal for executing machine learning tasks.
- Flexible Deployment Options VAIA can be:
 - Directly deployed in containers to enhance scalability, reduce latency, and bolster security, depending on infrastructure and operational needs
 - Deployed on top of a hypervisor layer to logically isolate hardware resources among different workloads in case of multitenancy.
- **Scalability** The infrastructure should be designed to scale smoothly, handling increased workloads over time without disruption.

Networking

An efficient network architecture is essential to ensure the responsiveness and reliability of VAIA application services. For seamless data transmission:

- **Bandwidth Planning** High Definition (HD) video streams generally require 2–5 Mbps per feed, while 4K streams can demand up to 25 Mbps, which must be considered during capacity planning.
- **Network Security** Employ firewalls, encryption, and Virtual Local Area Network (VLAN) segmentation to protect data and prevent unauthorized access.
- **Traffic Segregation** Isolate management, operational, and training data streams to ensure robust connectivity and prevent congestion. This can be achieved through physical separation or VLAN configuration.
- **VPN Connectivity** Reliable VPN connectivity is needed to ensure smooth integration of private and public cloud instances between inferencing and application, retraining services.

Storage integration

Efficiently managing the large volume of data generated by VAIA demands a carefully designed and scalable storage strategy.

- **High-Speed Storage** Non-Volatile Memory (NVMe) Solid State Drives (SSDs) are recommended to ensure fast data access and real-time processing.
- **Organized Data Management** Use structured storage and indexing methods to allow quick retrieval and analysis of recorded video content.

Sizing Considerations

To create a balanced and efficient deployment, system sizing should reflect expected operational demands:

- Al Workload Requirements Determine processing needs based on the number of simultaneous Al-analyzed video streams and enabled use-cases.
- Memory and Storage Provisioning Allocate enough memory and storage to avoid system bottlenecks.
- Deployment Models Assess configurations for small, medium, or large-scale deployments, depending on whether the use case involves a handful of HD feeds or complex, multi-location 4K processing.

Solution Validation Results

The solution has been validated across three distinct inference scales, as detailed in the table below:

	ThinkEdge Platform	GPU	CPU	RAM	Storage
Small	SE360 V2	1x Nvidia L4	16 cores	32 GB	1x 960 GB
Medium	SE360 V2	2x Nvidia L4	16 cores	64 GB	2x 960 GB
Large	SE455 V3	3x Nvidia L4	32 cores	128 GB	2x 960 GB

Integration is optimized for 1080p resolution video streams at 30 frames per second (FPS), ensuring highquality, real-time analytics. Further details about verification hardware platforms can be found in 'Appendix: Bill of Materials' section.

Validation Methodology

The Lenovo Validated Design (LVD) is verified through extensive testing in a controlled lab environment. Interoperability of solution components is validated in a controlled lab environment using commercially available far-edge devices such as cameras. Simultaneously, performance testing is conducted using simulators to emulate real-world workloads and stress conditions.

Figure 5 illustrates validation methodology, detailing the integration of data sources, the System Under Test (SUT), and the expected verification outputs.



Figure 5 - Validation Methodology

Summary of Validation Results

LVD for AI-Powered Industrial Workplace Safety has undergone comprehensive validation to ensure highperformance, real-time AI inferencing across various deployment scales. The solution was tested using Lenovo ThinkEdge SE360 V2 and SE455 V3 systems configured with NVIDIA L4 GPUs, simulating small, medium, and large deployment scenarios.

Each validation scenario assessed the solution's ability to process 1080p video streams at 30 FPS, confirming that each GPU reliably supports approximately 25 concurrent video streams, assuming adequate system resources. The validation process included not only performance benchmarking, but also installation, configuration, data integrity, and interoperability testing across edge and cloud components.

Key findings include:

- GPU as Primary Bottleneck: In all test cases, the GPU was the primary performance limiter.
- **Memory Impact**: A minimum of 64 GB Random Access Memory (RAM) is recommended for medium and large configurations to prevent memory-related bottlenecks.
- **Robustness Across Scales**: The solution scaled effectively across the three models, maintaining real-time analytics and consistent performance under load.

Complete platform configurations and detailed test outcomes are provided in Tables 3 and 4, with corresponding hardware specifications in the 'Appendix: Bill of Materials.'

Small Sizing m	Small Sizing model – SE360 V2 with 1x L4							
Detection &	Install	Config	Data	Interop	GPU	CPU	MEM	25 Cameras
Classification	Passed	Passed	Passed	Passed	High	Low	High	1080P, 30 FPS
Medium Sizing Model – SE360 V2 with 2x L4								
Detection &	Install	Config	Data	Interop	GPU	CPU	MEM	50 Cameras
Classification	Passed	Passed	Passed	Passed	High	Low	High	1080P, 30 FPS
Large Sizing Model – SE455 V3 with 3x L4								
Detection &	Install	Config	Data	Interop	GPU	CPU	MEM	78 Cameras
Classification	Passed	Passed	Passed	Passed	High	Low	Moderate	1080P, 30 FPS

Table 4 - Solution Validation Results

Solution Summary

The Lenovo Validated Design for AI-Powered Industrial Workplace Safety delivers a turnkey edge AI solution designed to enhance safety, reduce operational risk, and ensure regulatory compliance in industrial environments. Developed in partnership with Avathon, this solution combines Lenovo's rugged ThinkEdge SE360 V2 and SE455 V3 servers with Avathon's Visual AI Advisor (VAIA) platform to enable real-time detection of safety violations, PPE non-compliance, and hazardous conditions.

The VAIA platform inferencing runs advanced AI models directly at the edge, analyzing 1080p video feeds at 30 FPS from existing CCTV, drone, and mobile camera systems. The architecture further supports cloudbased model training together with edge-based inferencing, ensuring ultra-low latency, high reliability, and data privacy at scale.

Validated across small, medium, and large configurations, the solution scales efficiently based on customer needs. Each GPU supports approximately 25 video streams, and a minimum of 64 GB RAM is recommended for sustained performance at larger scales.

With Lenovo's XClarity and LOC-A for remote management and zero-touch provisioning, this modular and containerized solution enables rapid deployment, continuous learning, and centralized control—empowering safety teams to automate incident detection, streamline compliance, and drive safer, more efficient operations.

Appendix: Lenovo Bill of Materials

SE360 V2 BoM for Small

Part Number	Product Description	Qty
7DAMCTO2WW	ThinkEdge SE360 V2 for Ubuntu Server – 3y Warranty	1
BS56	ThinkEdge SE360 V2 Chassis	1
BS58	ThinkEdge SE360 V2 4x 1Gb, 2x 2.5Gb(TSN) I/O Module	1
BW2H	System Operational Temperature 5C to 40C / 41F to 104F	1
BFYE	Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
BS3X	ThinkEdge SE360 V2 Planar with Intel Xeon D-2775TE 16C 100W 2.0 GHz	1
B963	ThinkSystem 16GB TruDDR4 3200 MHz (2Rx8 1.2V) RDIMM	2
5977	Select Storage devices - no configured RAID required	1
AVV0	On Board SATA Software RAID Mode	1
BC4V	Non RAID NVMe	1
BS5M	ThinkEdge SE360 V2 M.2 Cabled Adapter Module	1
BQUJ	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	1
BS2C	ThinkSystem NVIDIA L4 24GB PCIe Gen4 Passive GPU	1
BS5J	ThinkEdge SE360 V2 Riser Assembly (PCIe Riser + 7mm Backplane)	1
BUGM	Internal AC Power Supply (230V/115V)	1
BUGU	ThinkEdge SE360 V2 AC Power Input Board	1
BW8U	ThinkEdge SE360 V2 500W 230V/115V Non-Hot Swap Power Supply	1
BS5W	ThinkEdge SE360 V2 Fan Assembly (Front to Rear)	1
BYBQ	XClarity Controller Managed	1
BOMK	Enable TPM 2.0	1
BRPJ	XCC Platinum	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

Part Number	Product Description	Qty
7DAMCTOBWW	ThinkEdge SE360 V2 for Ubuntu Server – 3y Warranty with Controlled GPU	1
BS56	ThinkEdge SE360 V2 Chassis	1
BS58	ThinkEdge SE360 V2 4x 1Gb, 2x 2.5Gb (TSN) I/O Module	1
BW2H	System Operational Temperature 5C to 40C / 41F to 104F	1
BFYE	Operating mode selection for: "Efficiency - Favoring Performance Mode"	1
BS3W	ThinkEdge SE360 V2 Planar with Intel Xeon D-2775TE 16C 100W 2.0 GHz	1
BQ37	ThinkSystem 32GB TruDDR5 4800MHz (2Rx8) RDIMM-A	2
5977	Select Storage devices - no configured RAID required	1
AVV0	On Board SATA Software RAID Mode	1
BC4V	Non RAID NVMe	1
BS5M	ThinkEdge SE360 V2 M.2 Cabled Adapter Module	1
BQUJ	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD (with Heatsink)	2
BS2C	ThinkSystem NVIDIA L4 24GB PCIe Gen4 Passive GPU	2
BS5J	ThinkEdge SE360 V2 Riser Assembly (PCIe Riser + 7mm Backplane)	1
BUGM	Internal AC Power Supply (230V/115V)	1
BUGU	ThinkEdge SE360 V2 AC Power Input Board	1
BW8U	ThinkEdge SE360 V2 500W 230V/115V Non-Hot Swap Power Supply	1
BS5W	ThinkEdge SE360 V2 Fan Assembly (Front to Rear)	1
BYBQ	XClarity Controller Managed	1
B0MK	Enable TPM 2.0	1
BRPJ	XCC Platinum	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

SE360 V2 BoM for Medium

SE455 V3 BOM for Large

Part Number	Product Description	Qty
7DBYCTOAWW	ThinkEdge SE455 V3 - 3Yr Warranty with Controlled GPU	1
BVTK	ThinkEdge SE455 V3 Chassis	1
BE4T	ThinkSystem Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-port OCP Ethernet	1
	Adapter	
BW2H	System Operational Temperature 5C to 40C / 41F to 104F	1

BFYA	Operating mode selection for: "Maximum Efficiency Mode"	1
BY8K	Acoustic Mode Disabled	1
BVTL	ThinkEdge SE455 V3 Motherboard	1
BY8X	ThinkEdge SE455 V3 AMD EPYC 8324PN 32C 130W 2.05GHz Processor	1
BW3L	ThinkEdge SE455 V3 1U Heatsink	1
BQ37	ThinkSystem 32GB TruDDR5 4800MHz (2Rx8) RDIMM-A	4
5977	Select Storage devices - no configured RAID required	1
BVUY	ThinkEdge SE455 V3 M.2 SATA/x4 NVMe Adapter with Carrier	1
BKSR	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BYM4	ThinkSystem 2.5" VA 960GB Mixed Use SATA 6Gb HS SSD v2	1
BVUV	ThinkEdge SE455 V3 4x2.5" SAS/SATA Backplane	1
BS2C	ThinkSystem NVIDIA L4 24GB PCIe Gen4 Passive GPU	3
BVUR	ThinkEdge SE455 V3 Riser1	1
BVUS	ThinkEdge SE455 V3 Riser2	1
BVV6	ThinkEdge SE455 V3 Intrusion Switch	1
BVTX	ThinkEdge SE455 V3 Standard EIA Bracket	1
BVUK	ThinkEdge SE455 V3 Power Distribution Board	1
BRH9	ThinkEdge 1100W 230V Titanium Hot-Swap Gen2 Power Supply	2
BVTP	ThinkEdge SE455 V3 Fan	5
BYBQ	XClarity Controller Managed	1
C1QS	SED Disabled	1
BPKR	TPM 2.0	1
BRPJ	XCC Platinum	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

Appendix: Abbreviations

Abbreviation	Meaning
AI	Artificial Intelligence
AMD	Advanced Micro Devices
API	Application Programming Interface
BSS	Business Support System
ССРА	California Consumer Privacy Act
CCTV	Closed-Circuit Television
CPU	Central Processing Unit
DC	Data Center
DVR	Digital Video Recorder
ERP	Enterprise Resource Planning
FPS	Frames Per Second
FQDN	Fully Qualified Domain Name
GDPR	General Data Protection Regulation
GigE	Gigabit Ethernet
GPU	Graphics Processing Unit
GUI	Graphical User Interface
HD	High Definition
HSE	Health, Safety, and Environment
НТТР	Hyper-Text Transfer Protocol
HTTPS	Hyper-Text Transfer Protocol Secure
HW	Hardware
I/O	Input/Output
loT	Internet of Things
JSON	JavaScript Object Notation
KPI	Key Performance Indicator
LOC-A	Lenovo Open Cloud Automation
LTS	Long-Term Support
LTSC	Long-Term Servicing Channel

LVD	Lenovo Validated Design
MES	Manufacturing Execution System
MQTT	MQ Telemetry Transport
NLP	Natural Language Processing
NVMe	Non-Volatile Memory
NVR	Network Video Recorder
nZTP	Near-Zero Touch Provisioning
OCP	Open Compute Project
OS	Operating System
OSS	Operations Support System
OSHA	Occupational Safety and Health Administration
PPE	Personal Protective Equipment
PTZ	Pan-Tilt-Zoom
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
REST	REpresentational State Transfer
RTSP	Real-Time Streaming Protocol
SFTP	Secure File Transfer Protocol
SSD	Solid State Drive
SSH	Secure Shell
SUT	System Under Test
TDP	Thermal Design Power
TLS	Transport Layer Security
TSN	Time Sensitive Network
UDP	User Datagram Protocol
UEFI	Unified Extensible Firmware Interface
VAIA	Visual AI Advisor
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPN	Virtual Private Network

VROC	Virtual RAID on CPU
WLAN	Wireless Local Area Network
WS	Web Service

Resources

#	Resources	Links
1	Avathon	Avathon Website
2	SE360 V2	Lenovo ThinkEdge SE360 V2 Server Product Guide > Lenovo Press
3	SE455 V3	Lenovo ThinkEdge SE455 V3 Server Product Guide > Lenovo Press
4	Lenovo XClarity	XClarity Systems Management
5	Lenovo Open Cloud Automation (LOC-A)	Lenovo Open Cloud Automation (LOC-A)
6	Visual AI Product Sheet	https://avathon.com/wp- content/uploads/2024/11/Visual-AI-Product- Sheet-A-VAA-PS-1162024-v1.0- 1.pdf?utm_source=chatgpt.com

Document history

[If there have been multiple versions of the Reference architecture document, then it makes sense to introduce version numbering and a history of changes to the document. This is why the version number is optional on the title page and page heading – it is not necessary for the first or even second version of a document.]

[Use a table to describe the list of changes as shown in the example below.]

Version 1.0April 2025First version includes Lenovo Edge Server SE360 V2 and SE455 V3
and Avathon's Visual Al Advisor (VAIA)

Trademarks and special notices

© Copyright Lenovo 2025.

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:

Advanced Micro Devices, AMD, AMD EPYC[™], Fire[™], and Xilinx 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®, Azure®, 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.

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

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

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

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

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

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

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

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

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