Lenovo Validated Design: AI-Powered Quality Inspection in Manufacturing

Last update: **15 March 2025**

Version 1.0

Trifork Vision AI and Lenovo Edge systems automate quality detection in manufacturing

Validated performance ensures ease of deployment, optimized AI workloads, and robust security

Scalable and Flexible Architecture, seamlessly integrates into manufacturing environments

Enhances quality control, automates inspections, ensures compliance, and enables real-time insights

Vanita Meyer

Sinan Atan

Table of Contents

Introduction

The Lenovo Validated Design (LVD) with Trifork document provides a comprehensive, pre-tested, and validated solution architecture for AI-powered quality management in manufacturing. This LVD outlines the integration of Lenovo ThinkEdge servers with Trifork Vision AI to deliver real-time defect detection and object classification, enabling manufacturers to enhance production efficiency, reduce defects, and ensure high product quality.

This document serves as a technical and business guide for organizations looking to deploy AI-driven vision-based quality control in their manufacturing processes. It details the business challenges, value proposition, solution components, deployment architecture, and validation methodology to help stakeholders understand and implement the solution effectively.

The LVD provides:

- A structured framework for AI-powered quality inspection in manufacturing.
- Reference architecture and validated sizing models to ensure seamless deployment.
- Performance benchmarks and interoperability testing to support industry standards.

By integrating AI, Edge, and IoT technologies, this LVD helps manufacturers:

- Optimize manufacturing QA investments by enhancing the effectiveness of quality assurance efforts
- Enhance product quality, efficiency, and customer satisfaction through automated defect detection
- Reduce quality control costs by minimizing defects and rework
- Minimize manual inspection efforts with AI-powered automation
- Ensure compliance with industry standards through consistent, high-accuracy inspections

Intended Audience

This document is intended for multiple stakeholders involved in the adoption of AI-driven smart manufacturing solutions, including:

- Manufacturing IT and Operations Teams – To understand how AI-powered vision-based quality control can be integrated into existing production environments.
- Enterprise Architects and Solution Engineers – To gain insights into the technical architecture, hardware specifications, and AI model integration for high-performance deployment.
- AI and IoT Solution Developers – To explore the software stack, model training, and inferencing capabilities of the solution.
- Business Leaders and Digital Transformation Executives – To evaluate the business impact, cost benefits of automating quality control processes.
- Channel Partners and System Integrators – To leverage a pre-validated AI solution that accelerates customer adoption while reducing implementation risks.

Scope of Work

This complete Vision AI based Quality Inspection solution as defined in this LVD is intended to be delivered in two phases. The breakdown of the components of each of those phases is outlined in the Architectural Overview and Operational Model sections. The validation efforts described throughout this document are limited to Phase 1, which focuses on the integration of edge devices with inferencing while ensuring management and reporting capabilities. Phase 2 expands the solution with new use cases, local model training, and fine-tuning at the edge, leveraging ThinkEdge SE455 V3 for AI retraining, data aggregation, and insights generation and presentation.

Business Problem and Business Value

Business Problem

The highly competitive nature of the manufacturing industry, combined with high customer expectations and tight margins requires industry players to meticulously manage product quality. Poor manufacturing output quality translates into declining customer satisfaction and trust, reduced revenues, and rising cost in the areas of returns processing, product rework and inspection intensity.

Traditional manual quality inspections are slow, labor-intensive, prone to human error, and costly due to high rates of defects, rework, and waste. Moreover, ensuring compliance with industry standards requires precise and consistent defect detection, something difficult to achieve with traditional quality control processes.

The rise of computer vision provides manufacturers with a powerful solution to address this challenge. Trifork Vision AI offering, powered by Lenovo enables companies to autonomously assess manufacturing quality and gain real-time insights, allowing them to proactively adjust production and quality processes before defects impact the end customer.

Business Value

The Trifork Vision AI solution, as outlined in this LVD, introduces a next-generation, AI-powered quality inspection system that revolutionizes manufacturing quality control. By leveraging computer vision, edge computing, and AI-driven analytics, it enables faster, more accurate, and automated defect detection, ensuring higher efficiency and product reliability.

Deploying cameras and edge computing devices in the manufacturing environment and integrating them with the organization’s enterprise network and systems enables:

- Instant visual capture of manufacturing production output
- Near real-time assessment of product quality gaps
- Accurate tracking of quality performance trends over time
- Actionable insights for manufacturing QA teams to enhance decision-making
- Continuous evaluation of QA investments and their impact
- Seamless knowledge sharing across facilities to standardize best practices

The integration of Vision AI into manufacturing QA efforts introduces a powerful new source of analytical data, driving measurable improvements across quality control processes. Key benefits include:

- Enhanced effectiveness of targeted QA initiatives
- Significant reduction in manual QA inspection efforts and costs
- Minimized waste and rework, improving operational efficiency
- Expanded coverage of product quality inspections
- Greater accuracy and efficiency in defect detection and assessments
- Higher production effectiveness, leading to increased profitability and customer satisfaction
- Consistent product quality across production lines and facilities

Finally, the combination of Lenovo technologies and Trifork Vision AI offering provides customers with the additional benefits of:

- **Scalability & Flexibility** – Pre-validated ThinkEdge servers ensure optimized AI workloads across various factory environments, allowing manufacturers to scale their AI capabilities efficiently.
- **Faster Time-to-Value** – Pre-tested and validated reference architecture streamlines deployment, reducing integration complexities and ensuring a quick return on investment (ROI).
- **Edge-to-Cloud AI Processing** – Supports real-time AI inferencing at the edge, reducing cloud-processing delays, optimizing latency, and enhancing data security with localized processing.
- **Compliance with Industry Standards** – AI-driven quality control ensures that manufacturers meet strict regulatory and compliance requirements with greater accuracy and reliability.

By adopting AI-powered vision-based quality inspections, manufacturers can move from relying on expensive, inaccurate and ineffective defect detection to high impact inspections that give QA teams the data they need to streamline the quality assurance process, unlocking higher operational efficiency, cost savings, and improved customer satisfaction.

Technology Overview

This section provides a comprehensive exploration of the technologies deployed and the requirements addressed in AI-powered quality management solution with Trifork and Lenovo.

Functional Requirements

The Trifork Vision AI solution is designed to empower manufacturers with advanced tools, capabilities, and insights that streamline quality assurance processes, reducing the time, effort, and costs associated with inspections. By delivering real-time, data-driven insights, it enables organizations to accurately measure the impact of QA investments and drive continuous improvements. This comprehensive solution addresses both the business and technical aspects of manufacturing quality, ensuring measurable enhancements in product consistency, operational efficiency, and overall business performance.

- **Problem Definition and Business Case Development:** Structured approach to identifying opportunities for improvement, defining the link between insights and operational changes, and development of a business case for the Vision AI journey and investment.
- **AI-Driven Computer Vision:** Leverages advanced camera technology to capture and digitize products in the manufacturing process. Interpret those images against defined output conditions and provide related insights to QA teams to inform operational changes aimed at enhancing efficiency and decision-making.
- **Seamless Enterprise Integration:** Enable closed-loop automation through integration to enterprise systems, deployment of robotics, optimization of workflows and delivery of operational efficiencies.
- **Scalable:** A flexible, AI-powered solution designed to fit seamlessly into any enterprise, ensuring smooth interaction between hardware and software layers.
- **Business-Centric Approach:** Develop a thoughtful approach to the initial deployment scope and the scale up and rollout strategy to optimize the cost / benefit of the Quality Inspection and related operational investments.

Non-functional Requirements

In addition to functional requirements, the following are the essential non-functional requirements considered during the solution's design phase to ensure it meets enterprise-grade standards:

- **Scalability:** The solution can scale both horizontally and vertically to accommodate various deployment considerations and future growth.
- **Fault Tolerance:** The system remains operational even if individual components fail, ensuring high availability.
- **Compact Physical Footprint:** The solution is designed for edge deployments without requiring a dedicated data center (DC).
- **Deployment Flexibility:** Supports multiple deployment models to adapt to different customer environments and infrastructure constraints.

- **High Performance:** Each component is optimized for speed and efficiency to prevent bottlenecks under heavy workloads.
- **Simplified Installation:** The deployment process is streamlined to minimize complexity and installation time.
- **Efficient Management & Operations:** The solution is designed for easy monitoring, maintenance, and operational efficiency.
- **Security & Data Protection:** Robust security measures safeguard customer data and infrastructure against threats and vulnerabilities.

Architectural Overview

The solution is deployed on Lenovo's ThinkEdge products for inferencing, with an optional AI Model Retraining capability. If AI Model Retraining is not deployed, the solution leverages a pre-trained model for inferencing, eliminating the need for retraining or fine-tuning in a specific environment.

The solution architecture is structured into two distinct planes, each serving a critical function:

Traffic Domain

The traffic domain handles the transmission of video recordings or images from production line cameras with minimal latency over a high-speed network. Data can be streamed live via Real-Time Streaming Protocol (RTSP) or captured on demand through a web service (WS).

Training and data lake functions of Phase 2 are also part of Traffic Domain. These establish a high-bandwidth, low-latency connection between the Inferencing and Training functions to enable fine-tuning of the solution for specific environments inclusive of:

- Continuous retraining and optimization using captured input data
- Adaptive improvements in model performance over time

AI Model Retraining at the Edge

- Phase 1 relies primarily on pre-trained models for inferencing
- Phase 2 would introduce on-device retraining using newly acquired data
- The 2x Nvidia L40S GPUs provide the computational power required for model fine-tuning, allowing the system to adapt to evolving manufacturing conditions.
- This capability enhances detection accuracy by continuously improving AI models based on newly acquired data.
- Retraining at the edge reduces cloud dependency, enhances data privacy, and improves response times.

Data Aggregation

- During Phase 2, a SE455 V3 would be deployed to collect and process high-volume data streams from multiple sources, including cameras, LiDAR, and sensors.
- By aggregating data locally, it minimizes latency and reduces the dependency on cloud infrastructure for real-time insights.
- This approach ensures efficient storage and organization of data, enabling better analysis and decision-making.

Management Domain

Operates on a high-priority, preferably dedicated network for system management, security, outbound quality monitoring and production line KPI analysis. It supports:

- Object-level observations via web service (WS)
- Batch processing for in-depth production belt analysis within System Observation modules
- System management through Secure Shell (SSH) for command-line control, along with remote desktop access and port forwarding for GUI-based operations
- A reporting module that retrieves solution KPIs via WS, providing actionable insights

Figure 1 illustrates the high-level deployment architecture of the proposed solution within this Lenovo Validated Design.

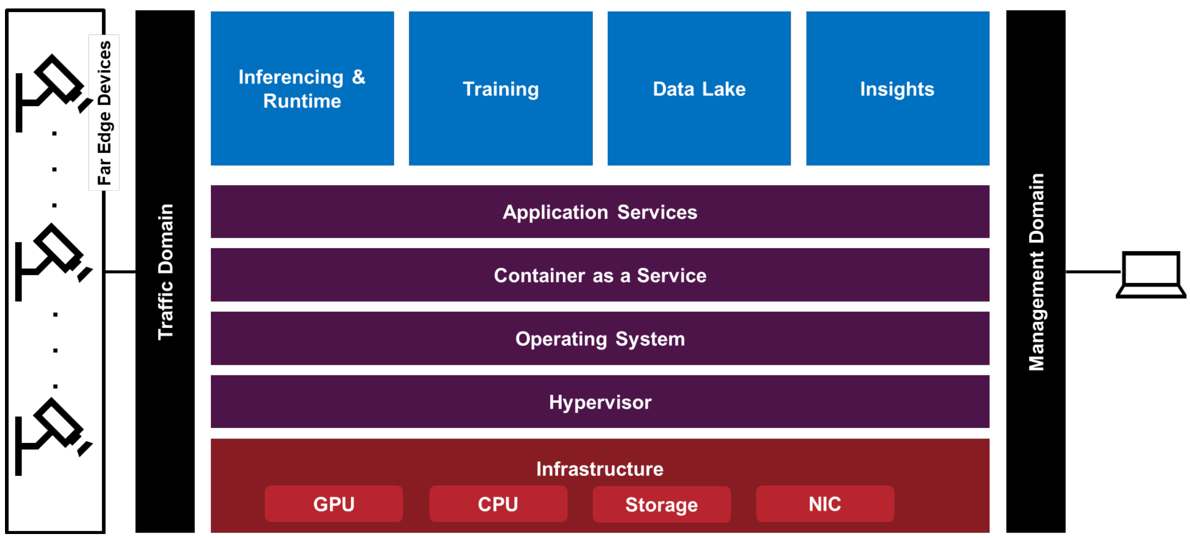

Figure 1 - Vision AI Driven Quality Inspection Deployment Architecture View

Key Features of the Architectural Diagrams:

- **Component Interactions**: The diagrams illustrate how AI components interact with hardware and software modules, including data capture, real-time processing, model inferencing, and reporting.
- **Data Flow**: The data pipeline is clearly outlined, from initial capture of product data to AI-driven defect or anomaly detection and quality assessment.
- **Infrastructure Layout**: The layout provides an overview of necessary infrastructure, including compute resources, storage, and network configurations.

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

Table 1 - LVD Technology Mapping

Layer	Technology	Version	Deployment Phase	Role
Far Edge (Cameras)	Basler Ace	ace 2 R a2A2448-23gcPRO	Phase 1	GigE PoE Color Camera 1080p @5 fps
Edge Infrastructure	ThinkEdge	SE360 V2 SE455 V3	Phase 1 (SE360 V2) Phase 2 (SE455 V3)	SE360 V2 for runtime & inferencing; SE455 V3 for training, data lake & insights
Hypervisor	KVM, OpenStack	TBD	TBD	Optional Layer, possible to validate further options
Operating System	Ubuntu 22.04	22.04	Phase 1	LTS version
CaaS	Docker	28.0.1	Phase 1	Container management (K8, etc.) as option for multi-modal deployments
Application Services	CUDA	12.2	Phase 1	Provides GPU acceleration for deep learning
	TensorRT	8.6.1.6		Optimizes AI inference performance on NVIDIA GPUs
	Nvidia Container Toolkit	2.13.0		Enables GPU support for containerized AI workloads
	Triton Inference Server	2.38.0		Scalable AI model deployment and inference optimization
	Astral UV	0.6.4		Image processing
Inferencing & Runtime	Trifork Vision AI	0.1.0	Phase 1	Real-time AI model execution for defect detection.
Model Training	Trifork Vision AI	0.1.0	Phase 2	AI model retraining for improved accuracy and adaptability
Data Lake	Trifork Vision AI	0.1.0	Phase 2	Centralized storage for analytics and model training
Insights/dashboard	Trifork Vision AI	0.1.0	Phase 2	Data driven decision making

Solution Roadmap

Phase 1: Inferencing with Pre-Trained Model

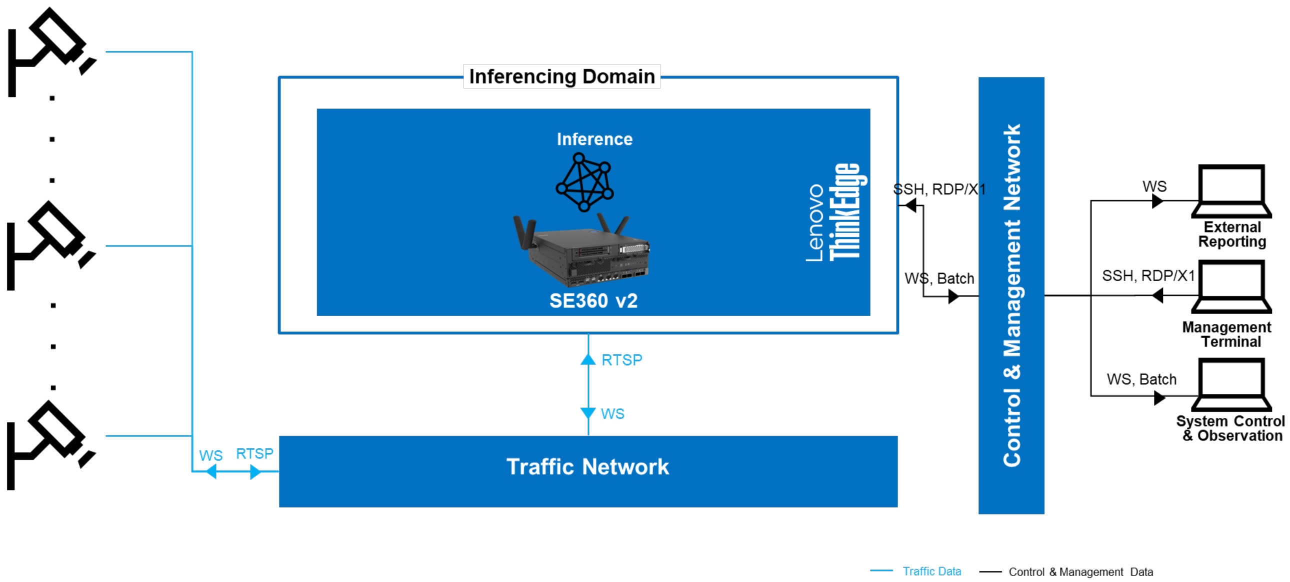

Figure 1 - Phase 1 Deployment Architecture Overview

Phase 2 : Inferencing with Local Training & Fine-Tuning

Phase 2 aims to expand the LVD by incorporating Training and Insights domains with the introduction of the ThinkEdge SE455 V3. The SE455 V3 is designed to handle AI model training, data aggregation/lake, and insights functions, enhancing overall system capabilities.

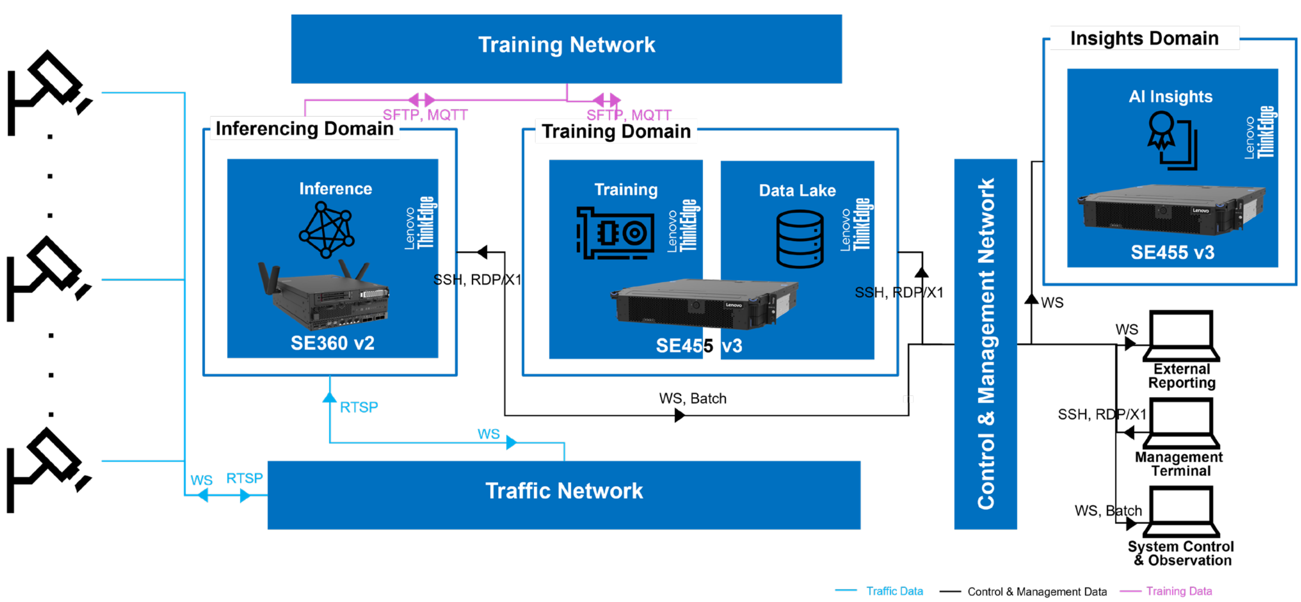

Figure 2 - Phase 2 Deployment Architecture Overview

Component Model

Figure 2 illustrates the infrastructure supporting the quality inspection solution, which leverages computer vision and machine learning. Designed for manufacturing environments, the solution analyzes geometric properties, detects anomalies, defects, and inconsistencies in production lines, and enables automated quality assessments to enhance efficiency and accuracy.

This solution supports the integration multiple AI and computer vision models, such as defect detection and classification, enabling real-time processing of real-time data from cameras, LiDARs, lasers, and other sensors as needed. This highly flexible infrastructure deployed across various manufacturing environments, adapting to specific operational needs, including:

- GPU acceleration for high-performance computing
- Containerization for scalability and streamlined deployment
- On-premises hardware for environments requiring local processing

Additionally, the solution allows seamless output connections to various actuation and control systems, including:

- Actuators and Programmable Logic Controllers (PLCs) for automated responses
- User interfaces for real-time monitoring and decision-making
- Indicator lights to signal detected product nonconformities and trigger immediate corrective actions

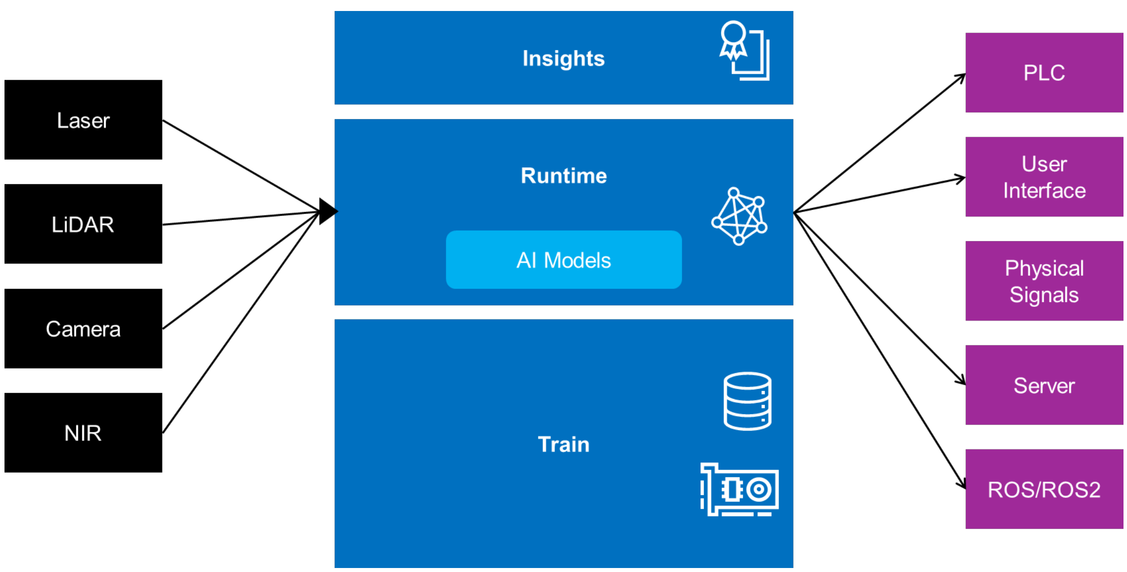

Figure 4 - Quality Inspection Application Architecture

Trifork Vision AI Application Services

The architecture of the quality inspection solution is intended to support integration with existing manufacturing workflows. The solution comprises multiple components that interact to capture, process, analyze, and report inspection results via three components.

Insights

The insights component allows for custom development and integration of capabilities for reporting and analyzing business metrics and KPIs, translating inspection data into actionable insights. Customers can further customize this component to include tailored insights, enabling specific reporting needs and deeper analytics based on their unique manufacturing requirements.

Runtime

Delivers the operational workflow for data processing and monitoring the solution. It includes:

- **Operation Monitoring**: Monitors the health of devices and processes.
- **Workflow**: Customizable inspection procedures, enabling adaptation to different production setups and requirements.
- **Model Inference**: Scalable and flexible model serving to handle multiple models for developing quality requirements and future needs.

Model Training

Provides the capability to improve and develop models for future quality assessment needs and requirements (phase 2).

Table 2 - Component Relationships

Input Hardware (Cameras, etc.)	Various devices (e.g., cameras, LiDARs, lasers) capture and relay information about products or materials on the production line. Positioning and configuration are subject to use-case requirements.
Edge Processing Unit	Edge computing devices process input data to reduce latency and enhance real-time defect detection. Actual model inference is performed with the Model Inference Server.
Model Inference Server	A service that runs computer vision and deep learning models for quality assessment, supporting scalable deployment and flexible handling of inference requests.
Data Storage & Logging	Integrations to store inspection data, metadata, and model outputs for e.g. audits, further analysis, or compliance.
Production Line Integration	Interfaces with e.g. manufacturing execution systems (MES), PLCs, or other controllers to trigger actions, notifications, or corrective measures based on inspection outcomes.
Operation Monitoring	Observes the state of both the hardware (e.g., edge device health) and the software (e.g. resource utilization), allowing intervention when it might be required.

Insights	Offers data-driven insights, and KPI tracking to help decision-makers evaluate the effectiveness of the inspection process and respond appropriately.

By leveraging this quality inspection solution, organizations may reduce defective outputs, improve production workflows, and support overall quality standards while gaining valuable insights for ongoing process optimization.

ThinkEdge Platform Services

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** – Get an integrated, cost-effective, and resilient edge solutions that are easy to deploy, simple to operate, and maximize edge workload performance.
- **Improve Data Management and Access** – Better data management with the widest choice of products, services, and most importantly—guidance, available for advanced hybrid infrastructure.

Lenovo's ThinkEdge SE360 V2 and SE455 V3 are the recommended platforms for integrating Trifork Vision AI Quality Management solution.

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, 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 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

Table 3 - ThinkEdge SE360 V2 Characteristics

Form Factor	2U Height, Width: 212mm, Depth: 317.5mm
Processor	Intel® Xeon® D processor, up to 20 cores, up to 125W 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 (RAID 1) + 2 x M.2 Data drives internally
Memory	4 x DDR4 memory slots; Maximum 256GB
Expansion Slots	1 x PCIe G4 x16 HHHL + 1 x PCIe G4 x 16 FHHL
GPUs	Nvidia A2 and L4, Qualcomm Cloud AI 100, Intel Flex 140
Network Interface	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)
Front panel & I/O	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 Port (video) / USB3.0 auto switch), BMC serial
HBA/RAID Support	RAID 0, 1, 5, 10 (with Intel VROC 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)
Operating temperature	-20^{o}C to 65^{o}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

Lenovo's ThinkEdge SE455 V3 achieves most AI performance with best power efficiency. It is considered for Phase 2 of the LVD scope for GPU intensive training purposes, data lake and reporting capabilities. Below is the key characteristics of SE455 V3:

- **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 OCP achieving 2x more strage 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 UEFI settings, larger CPU heatsink
- **Energy Efficient** – 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** – Upt to 6x single-width GPUs or 2x double-width GPUs supported with upt 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

Table 4 - ThinkEdge SE455 V3 Characteristics

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
Memory	6x DDR5 memory slots, up to 576GB using 96GB RDIMM (128GB DIMM in plan) Common options with ThinkSystem

Expansion Slots	2x PCIe Gen5 x16 (Riser2 optional) 4x PCIe 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 Management	Lenovo XClarity Controller (AST2600) DC-SCM 2.0

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 LTSC
Limited Warranty	3-year customer replaceable unit and onsite service, next business day 9x5 Optional service upgrades

Operational Model

Lenovo Xclarity & LOC-A provides holistic dashboards to manage IT systems and workloads with advanced automation capabilities. It is as well possible to integrate those with open APIs to northbound systems to federate operation & maintenance capabilities.

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 web-based REST APIs for ease of inter-operability.
- **XClarity Essentials** – It ensures collection of one-to-one management scripting tools as ONECLI, UpdateXpress and Bootable Media Creator.
- **XClarity Administrator** – Management platform for ThinkEdge deployments scaling up to 1000 nodes. It exposes GUI capabilities for 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** – Energy Manager provides:
    - 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 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)

LOCA 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, late-binding 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 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:
    - Bare-metal server provisioning and onboarding
    - Orchestration of edge-node deployments
    - Edge cluster and node instance creation (including OS deployment)
    - Access to a smart naming convention engine for hostname, 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 OSS/BSS platforms. This API enhances integration capabilities beyond the plugin mechanism, accelerating new feature enablement across platforms

Deployment Considerations

Successful deployment of Trifork Vision AI requires a well-structured approach to ensure seamless integration with existing IT infrastructure. Key considerations include system management, server and compute resources, networking, and storage. This section outlines the critical deployment factors, including:

- **Performance optimization** to ensure real-time processing efficiency
- **Security measures** to protect data and infrastructure
- **Scalability** to accommodate future growth and evolving business needs

Additionally, it details how this LVD supports these capabilities, ensuring a robust and reliable solution. A high-level overview of the IT environment requirements necessary for implementing Trifork Vision AI is also provided to guide deployment planning.

Systems Management

A structured and efficient systems management strategy is crucial for maintaining Trifork Vision AI and ensuring continuous operations. To optimize system reliability and security, organizations should focus on:

- **Monitoring & Performance Tracking** – Connect Trifork Vision AI health and status information to your existing monitoring solutions, or use the default NetData integration, to track system health, detect potential issues early and capture detailed operational logs.
- **Automated Updates & Security Patching** – Deploy automated management tools to keep the software environment up to date with the latest security patches and feature updates, minimizing vulnerabilities and downtime.
- **Remote Administration** – Lenovo XClarity offers remote management capabilities, enhancing operational oversight for ThinkEdge-based deployments.

Server / Compute Nodes

The core component of Trifork Vision AI Runtime (Run) is engineered to support demanding AI applications, offering flexibility in deployment across centralized server hardware or decentralized edge hardware, depending on specific use cases and network requirements.

For optimal performance, deployments should be built around robust compute nodes capable of handling real-time video processing. Key deployment considerations include:

- **Dedicated AI Hardware** – Lenovo ThinkEdge servers equipped with high-performance CPUs, GPUs, and ample memory ensure efficient execution of machine learning workloads.
- **Flexible Deployment** – Trifork Vision AI can be deployed:
    - Directly on dedicated ThinkEdge servers running Ubuntu
    - In a containerized environment to optimize scalability, latency, and security, based on available hardware resources and operational constraints.

- **Scalability** – The infrastructure should be designed to seamlessly handle increasing workloads over time. It should be proactively planned to accommodate growing processing demands, ensuring smooth, efficient, and uninterrupted operation as AI/ML workloads scale.

Networking

A well-designed network is fundamental to ensuring the responsiveness and efficiency of Trifork Vision AI applications. To maintain seamless data flow:

- **Bandwidth Planning –** Standard HD cameras typically consume 2–5 Mbps per stream. 4K Video Feeds may require up to 25 Mbps, significantly impacting network capacity.
- **Security Best Practices** – Implementation of firewalls, data encryption, and VLAN segmentation ensures safeguarding network traffic and prevent unauthorized access.

The segmentation of management, traffic and training domain data streams ensures robust network connectivity avoiding cross-domain congestion. Physical separation of network domains or VLAN segmentation is possible for this purpose.

Storage Integration

Handling the vast amounts of data generated by Trifork Vision AI requires a carefully structured storage solution:

- **Fast & Efficient Storage** – NVMe SSDs are recommended for high-speed data access and processing.
- **Scalable & Flexible Storage Options** – On Premise or Cloud-Based Storage is possible for efficient long-term data retention and easy scalability as data volumes grow.
- **Organized Data Management** – Indexing and Structured Storage enable rapid retrieval and analysis of recorded video data, ensuring quick access to critical information.

Sizing Considerations

To ensure a well-balanced deployment, system sizing should be based on expected workloads:

- **Compute & AI Processing Needs** – Workloads vary depending on the number of AI-driven video streams being processed simultaneously.
- **Memory & Storage Allocation** – Sufficient memory and storage resources must be provisioned to avoid performance bottlenecks.
- **Deployment Models** – Organizations should evaluate small, medium, and large-scale configurations based on their unique use cases, whether managing a few HD streams or processing extensive multi-location 4K feeds.

By designing a deployment strategy that aligns with these key factors, businesses can optimize the performance, scalability, and reliability of their Trifork Vision AI implementation.

Solution Validation Results

The solution has been validated for three different inferencing scales in Phase 1, as outlined in the table below:

Table 5 - Validation Platforms

	ThinkEdge Platform	GPU	CPU	RAM	Storage
Small	SE360 V2	1x Nvidia A2	8 cores	16 GB	2x 960 GB
Medium	SE360 V2	1x Nvidia L4	16 cores	32 GB	2x 960 GB
Large	SE360 V2	2x Nvidia L4	16 cores	32 GB	2x 960 GB

Integration is done with 1080p resolution using 5 FPS video streams.

Verification Methodology

LVD verification is done in two phases where interoperability of the solution components is validated in the lab environment with the use of commercially available far edge devices (cameras) and performance testing is done with the help of simulators. Figure 4 provides a brief overview of Phase 1 verification methodology in case of data sources, System Under Test (SUT) and verification outputs.

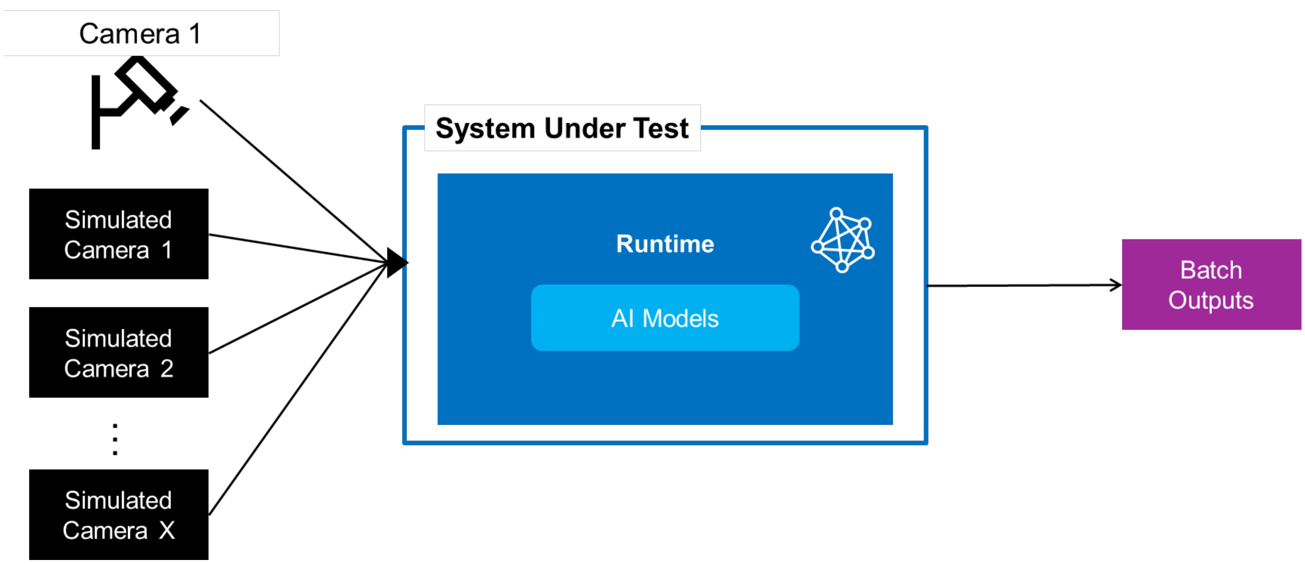

Figure 5 - Phase 1 Validation Methodology

Summary of Validation Results

The validation results for the LVD document assess the performance of different sizing models using SE360 V2 with various configurations, including A2 and L4 accelerators. These models are tested under both Basic and Advanced vision models to determine their suitability for object classification and defect detection tasks.

Basic vs. Advanced Vision Models

- **Basic Vision Models** are well-suited for simpler tasks, such as classifying distinct objects or detecting general defects. They typically use bounding boxes to highlight areas of concern but may lack the granularity needed for highly detailed inspections. Typically, identifying the color of products or detecting major packaging defects falls within the capabilities of a basic model.
- **Advanced Vision Models** are designed for more complex analysis. They are ideal for inspecting intricate products, where defect detection requires high accuracy. These models leverage precise position detection, enabling detailed segmentation (e.g., exact masking of defects rather than bounding boxes). This level of precision is crucial for applications where minor imperfections can have significant consequences.

Table 6 – Validation Results

Small Sizing model – SE360 V2 with A2								
Basic	Install Passed	Config Passed	Data Passed	Interop Passed	GPU High	CPU Moderate	MEM Moderate	140 Cameras 1080P, 5 FPS
Advanced	Install Passed	Config Passed	Data Passed	Interop Passed	GPU High	CPU Low	MEM Low	2 Cameras 1080P, 5 FPS
Medium Sizing Model – SE360 V2 with 1 x L4								
Basic	Install Passed	Config Passed	Data Passed	Interop Passed	GPU High	CPU Low	MEM High	192 Cameras 1080P, 5 FPS
Advanced	Install Passed	Config Passed	Data Passed	Interop Passed	GPU High	CPU Low	MEM Low	4 Cameras 1080P, 5 FPS
Large Sizing Model – SE360 V2 with 2 x L4								
Basic	Install Passed	Config Passed	Data Passed	Interop Passed	GPU High	CPU Moderate	MEM High	260 Cameras 1080P, 5 FPS
Advanced	Install Passed	Config Passed	Data Passed	Interop Passed	GPU High	CPU Low	MEM Moderate	6 Cameras 1080P, 5 FPS

Three different hardware configurations were evaluated:

- **Small Sizing Model – SE360 V2 with A2**:
  - **Basic Model**: Passed all tests and supported **140 cameras (1080P, 5 FPS)**. GPU usage was high, while CPU and memory utilization were moderate.
  - **Advanced Model**: Passed all tests, supported only **2 cameras (1080P, 5 FPS)** due to the higher computational demands of the advanced model.
- **Medium Sizing Model – SE360 V2 with 1 × L4**:
  - **Basic Model**: Passed all tests, supporting **192 cameras (1080P, 5 FPS)**. GPU was heavily utilized, CPU usage was low, and memory usage was high.
  - **Advanced Model**: Supported **4 cameras (1080P, 5 FPS)**, demonstrating that the advanced model requires significantly more processing power.
- **Large Sizing Model – SE360 V2 with 2 × L4**:

- **Basic Model**: Successfully validated, handling **260 cameras (1080P, 5 FPS)** with moderate CPU and high memory utilization.
- **Advanced Model**: Passed with **6 cameras (1080P, 5 FPS)**, showing an incremental increase in camera capacity compared to the medium model.

Key Insights:

- **Basic models** scale well and can support a high number of cameras, making them ideal for general-purpose monitoring and object classification.
- **Advanced models** require significantly more resources but provide superior accuracy and precision, which is essential for critical defect detection applications.
- **Scaling hardware with more powerful GPUs (L4 vs. A2) significantly improves capacity**, especially for advanced models, allowing for a higher number of supported cameras.

The choice between basic and advanced vision models depends on the application: large-scale monitoring benefits from basic models, whereas high-precision tasks require advanced models.

Solution Summary

The Lenovo Validated Design (LVD) for Trifork Vision AI Quality Inspection for Manufacturing solution delivers powerful capabilities that streamline and accelerate the process of identifying product quality issues and enables manufacturers to quickly assess the tangible impact of operational changes on product quality, costs, and customer satisfaction.

By integrating Lenovo ThinkEdge SE360 V2 and SE455 V3 servers into Trifork Vision AI solution, organizations can benefit from real-time inferencing, automated quality inspections, and valuable insights at scale and at speed.

Key highlights of this solution include:

- **Instant Visual Capture of Production Output**
  Rapid and automated imaging of manufactured products to ensure comprehensive quality tracking.
- **Near Real-Time Quality Assessment**
  AI-driven analysis identifies product quality gaps almost immediately, enabling quick corrective actions.
- **Continuous Monitoring of Quality Performance Trends**
  Track changes in manufacturing quality over time, providing accurate insights into process improvements or inefficiencies.
- **Actionable Insights for QA Teams**
  Delivers real-time, data-driven quality insights to manufacturing QA teams, enabling proactive decision-making.
- **Impact Measurement of QA Investments**
  Quantifies the effectiveness of quality assurance initiatives, helping manufacturers assess ROI and optimize resources.
- **Enterprise-Wide Knowledge Sharing**
  Facilitates the standardization of best practices by sharing quality insights across multiple facilities.

Powered by these technologies:

- **AI-Driven Defect Detection & Object Classification** – Enhances accuracy, reduces errors, and minimizes production waste.
- **Optimized Edge AI Performance** – ThinkEdge servers ensure seamless deployment, efficient AI inferencing, and scalability.
- **Validated Performance & Interoperability** – Pre-tested configurations support flexible deployment models for varying workloads.
- **Scalable Architecture** – Supports a broad range of manufacturing environments, from small-scale deployments to large, high-performance implementations.

- **Reduced Operational Costs & Increased Efficiency** – Automation accelerates production, lowers rework rates, and improves resource utilization.
- **Edge-to-Cloud AI Processing** – Enables real-time insights while reducing latency and enhancing security.

The solution validation results confirm robust performance across multiple deployment models, showcasing the ability to scale from small to large configurations. With a structured roadmap for initial inferencing and future AI model retraining at the edge, this LVD provides manufacturers with a future-ready framework for AI-driven quality management.

By adopting this solution, organizations can streamline operations, achieve higher product consistency, meet compliance standards, and drive digital transformation in manufacturing.

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
BS3W	ThinkEdge SE360 V2 Planar with Intel Xeon D-2733NT 8C 80W 2.1 GHz	1
B963	ThinkSystem 16GB TruDDR4 3200 MHz (2Rx8 1.2V) RDIMM	1
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
BP05	ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU	1
BS5J	ThinkEdge SE360 V2 Riser Assembly (PCIe Riser + 7mm Backplane)	1
BUGM	Internal AC Power Supply (230V/115V)	1
BUGU	ThinkEdge SE360 V2 AC Power Input Board	1
BW8U	ThinkEdge SE360 V2 500W 230V/115V Non-Hot Swap Power Supply	1
BS5W	ThinkEdge SE360 V2 Fan Assembly (Front to Rear)	1
BYBQ	XClarity Controller Managed	1
B0MK	Enable TPM 2.0	1
BRPJ	XCC Platinum	1
7S0XCTO5WW	XClarity Controller Platin-FOD	1
SBCV	Lenovo XClarity XCC2 Platinum Upgrade (FOD)	1

SE360 V2 BoM for Medium

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
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)	2
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
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 Large

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
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)	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

SE455 V3 BOM for Training

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 Adapter	1
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	2
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
BYFH	ThinkSystem NVIDIA L40S 48GB PCIe Gen4 Passive GPU	2
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
DC	Data Centre
ERP	Enterprise Resource Planning
FQDN	Fully Qualified Domain Name
GigE	Gigabit Ethernet
GUI	Graphical User Interface
KPI	Key Performance Indicator
LOC-A	Lenovo Open Cloud Automation
LVD	Lenovo Validated Design
nZTP	Near-Zero Touch Provisioning
PoE	Power over Ethernet
ROI	Return on Investment
RTSP	Real-Time Streaming Protocol
SFTP	Secure File Transfer Protocol
SSH	Secure Shell
WS	Web Service

Resources

Resources	Links
Trifork	Trifork website
SE360 V2	Lenovo ThinkEdge SE360 V2 Server Product Guide > Lenovo Press
SE455 V3	Lenovo ThinkEdge 455 V3 Server Product Guide > Lenovo Press
XClarity	Systems Management
Lenovo Open Cloud Automation (LOC-A)	Lenovo Open Cloud Automation
Netdata	https://www.netdata.cloud/

Document history

Version 1.0	March 2025	AI-Powered Quality Inspection with basic and advanced models of Defect Detection and Object Classification use cases. Validated on the SE360 V2 for 3 sizing models.

Trademarks and special notices

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.