# Lenovo Validated Design for Retail: Simplified Deployment and Management

Last update: **10 October 2024**

Version 1.0

**LOC-A automates edge infrastructure deployment and lifecycle management across sites**

**Simplify application deployment and management with Intel® Tiber™ Edge Platform**

**Viana™ by meldCX® uses AI analytics to enhance retail operations, offering insights on customer behavior and engagement.**

**Lenovo ThinkEdge Systems deliver reliable edge computing performance for AI tasks**

**Vanita Meyer**

**Chris Arenas**

# Table of Contents

# Introduction

Edge computing represents a transformative approach to data processing that is rapidly gaining momentum. By processing data closer to its source rather than relying on distant cloud services, this methodology offers significant advantages in scenarios requiring immediate data analysis or minimal latency. According to Gartner, by 2025, enterprises are expected to manage over 75% of their data outside traditional data centers or cloud infrastructures, driving substantial growth in the edge computing market. To harness real-time insights and avoid the delays associated with cloud processing, organizations are increasingly positioning workloads near the points where data is generated.

In the retail sector, vast amounts of data are continuously generated from sources such as point-of-sale systems, customer traffic monitoring, digital signage engagement tracking, and security cameras. This makes retail environments ideal for edge computing, where the processing of data locally can lead to faster decision-making, enhanced customer experiences, and improved operational efficiency.

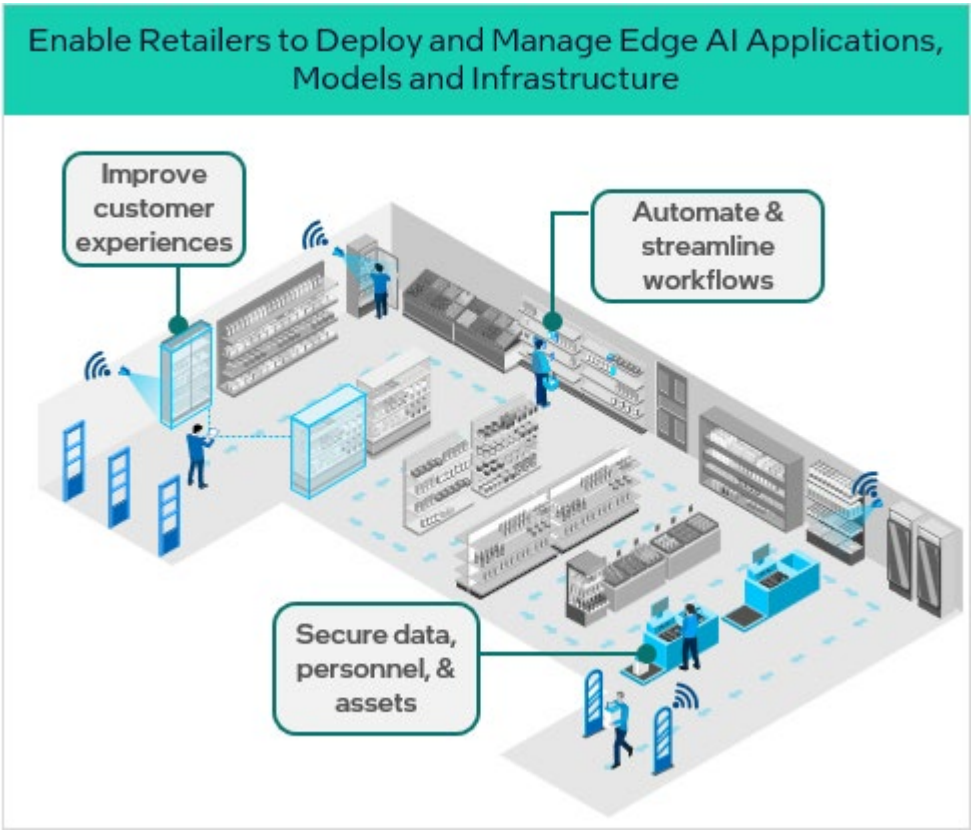


***Figure 1 - Retail Solutions Overview***

# Document Objective/Purpose

This document outlines Lenovo Validated Design (LVD) and Viana™ solutions for retail, leveraging the Edge Platform and Lenovo lifecycle management tools to simplify deployment and management. The goal of the Lenovo design and validation process is to confirm the performance, interoperability, and seamless workflow between systems and applications in a retail edge environment.

Viana (short for vision analytics) by meldCX® is an advanced platform designed to enhance retail operations through edge computing and AI-driven applications. It leverages the Intel® Tiber™ Edge Platform (Edge Platform) and Lenovo lifecycle management tools (Lenovo Cloud Automation and XClarity Controller) to simplify the deployment and management of retail edge applications, AI models and infrastructure. The solution enables near real-time data processing for tasks such as customer engagement, signage content effectiveness, and audience measurement, while providing valuable insights into the customer journey. By integrating seamlessly with existing systems, Viana enhances store operations, improves customer experiences, reduces losses, and helps retailers quickly implement new technologies for smarter, more efficient retail environments that can add value and unlock opportunities for added revenue.

**Intended Audience**

This document is designed for IT professionals, technical architects, sales engineers, channel partners, and consultants. Its purpose is to guide the planning, design, and implementation of retail edge solutions, offering platforms and tools that streamline the development, deployment, and management of these solutions. By providing clear frameworks and best practices, it aims to simplify complex processes and enhance operational efficiency across retail environments.

# Business challenges and business value

## Business challenges

- High-cost new technology integration
- Siloed fragmented applications
- Complex and dispersed infrastructure

As retailers face rising operational costs and evolving customer expectations, they are increasingly turning to edge computing technologies to remain competitive. However, the rapid proliferation of edge applications, combined with fragmented application silos, introduces significant complexity and technical challenges that hinder innovation. This complexity can slow the modernization of retail stores and complicate the integration of new technologies aimed at delivering operational efficiencies and significantly better customer experiences.

## Business value

Lenovo Validated Design (LVD) with Viana, powered by Intel Tiber Edge Platform and Lenovo lifecycle management tools, address these challenges by simplifying the deployment, orchestration, and management of edge applications, AI models, and infrastructure. These solutions improve manageability, speed, scalability, and resiliency at the edge, enabling smarter retail services and operational efficiency. Retailers can quickly evaluate and deploy new applications that provide valuable insights into customer journeys, opening new revenue streams, enhancing customer experiences, improving safety and security, and reducing losses—all while efficiently managing their infrastructure and applications.

### Benefits of the Integrated Solution

This solution offers several advantages for businesses adopting retail edge computing:

- **Accelerated Deployment & Scalability**: The Viana + Edge Platform provides a pre-validated framework for faster deployment, allowing businesses to scale seamlessly and support multiple AI models in parallel as needs grow.
- **Real-Time Processing, Reduced Latency & Flexibility**: By processing data closer to the source, the platform reduces latency for real-time decision-making in applications like customer engagement. It is also compatible with a wide range of Lenovo hardware and software, offering flexible deployment options.
- **Interoperability & Advanced AI Integration**: The platform integrates with Intel partners, like Hugging Face, supporting advanced AI and analytics via OpenVINO, enhancing functionality and customer experiences.
- **Enhanced Security & Data Privacy**: Viana's Zero Trust security model, combined with IAM, encryption, and a hardened OS, ensures data privacy. Captured data is anonymized, encrypted, and stored temporarily, with devices secured in Kiosk Mode.

- **Cost & Operational Efficiency**: The integration optimizes system performance and consolidates tools to reduce costs. Real-time automation boosts productivity, minimizes downtime, and maximizes resource use without over-investing in infrastructure.
- **Validated Solution & Resiliency**: The platform, validated with Lenovo lifecycle tools and Viana, ensures reliable deployments. It operates continuously, even during cloud connectivity disruptions.
- **Near-Zero Touch Provisioning & Simplified Infrastructure Management**: Lenovo Open Cloud Automation (LOC-A) automates deployments across multiple locations, while Lenovo XClarity Administrator (LXCA) simplifies provisioning, monitoring, and updates in hybrid environments.

By offering a robust, scalable, and secure framework for retail edge, this integrated solution enables businesses to accelerate innovation, optimize operations, and unlock new opportunities across retail environments.

# Test & Validation: Scope, Methodology, Results

The validation scope focuses on assessing the efficiency and performance of edge infrastructure deployments, including onboarding third-party applications, deploying operating systems, and ensuring seamless application compatibility with Edge Platform and Viana. It evaluates deployment speed, simplicity, and the effectiveness of automated processes. Additionally, it ensures reliable connectivity, optimized resource use, stable network performance, and provides guidelines for sizing and maintaining performance benchmarks within set guardrails.

## Validation Scope

- The objective is to test and validate the performance and interoperability of Lenovo ThinkEdge SE350 V2, SE360 V2, and SE450 servers with Viana, Edge Platform, and LOC-A with XCC and XClarity Administrator.
- This validation includes ensuring successful provisioning, lifecycle management, and the functioning of containerized AI retail inferencing workloads on the Intel CPU.

### Key Components:

- **Hardware**: Lenovo ThinkEdge SE350 V2, SE360 V2, SE450.
- **Software**: Viana (v3.6.1) by meldCX for AI inferencing and data analysis, Edge Platform (v24.08) for container orchestration, LOC-A (v2.8) for provisioning and lifecycle management.
- **Additional Setup**: An additional SE450 will act as a faux IP camera, generating RTSP streams mimicking the data from 1 – n IP cameras.

## Validation Methodology

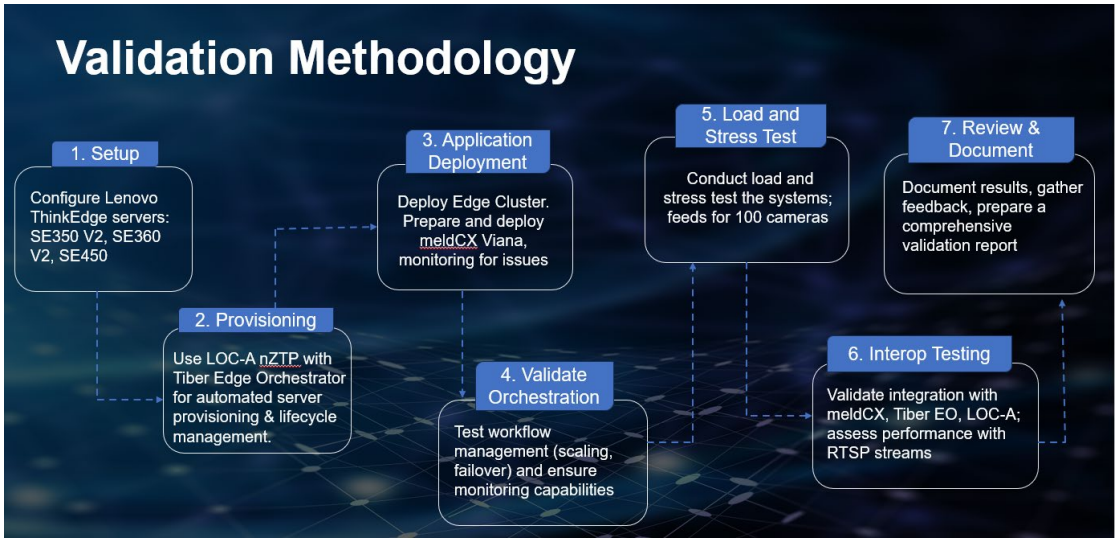


***Figure 2 – Validation methodology***

## Test Environment Setup

**Hardware**:

- **Lenovo ThinkEdge SE350 V2, SE360 V2, SE450**: Detail the configurations (memory, storage, network setup).
- **SE450 as a Faux IP Camera**: This server will simulate 1 to 100 IP camera streams by bombarding the SE350 V2 and SE360 V2 with RTSP feeds.

**Software Stack**:

- **Viana**: Installed on the ThinkEdge servers to perform retail AI inferencing such as people detection and data aggregation.
- **Edge Platform**: Manages the orchestration of containerized applications across the platform.
- **LOC-A + XCC + LXCA**: Handles provisioning, lifecycle management, and NZTP to automate server setup and configuration.

**Network Configuration**:

- All devices are on the same network segment, ensuring seamless communication and interaction.
- The lab environment guarantees a minimum network throughput of **1Gbps**.

# Test Cases

- **Provisioning via LOC-A:** Validate that the ThinkEdge servers are provisioned seamlessly using LOC-A's NZTP capabilities.
- **Application Orchestration**: Ensure Edge Platform can manage and orchestrate Viana containers across edge node clusters made of Lenovo ThinkEdge SE350 V2, SE360 V2, and SE450 systems.
- **RTSP Stream Handling**: Test the ability of the ThinkEdge servers to process video streams from the faux IP camera SE450, ranging from 1 to n streams.
- **Functional Testing of Viana**: Validate core functions, including AI inferencing tasks like people detection and retail analytics using simulated video streams.
- **Performance and Load Testing**: Measure the system's performance and load-handling capability as stream count increases, monitoring CPU/memory usage and throughput.
- **Scalability Testing**: Assess the scalability of Edge Platform in managing multiple containers and AI inference tasks while maintaining system performance and stability.

## Data Capture and Analysis

- **Monitoring Tools**: Utilize monitoring tools (e.g., Prometheus, Grafana) to capture metrics on CPU, memory usage, and network throughput during testing.

- **Captured Metrics**: Record data on CPU/memory load, network traffic, and AI inference accuracy for varying stream loads for 100 cameras.
- **Anomaly Detection**: Monitor for system errors, application crashes, or network congestion during the tests.

### Resiliency Testing

- **Failover and Recovery**: Test how the system responds to failure conditions, such as network disconnections or high system load.
- **Handling Overload Conditions**: Overload the system by increasing streams beyond its expected limit and measure its behavior under stress (e.g., system throttling or degradation).

### Reporting and Analysis

- Server performance during a sustained load, peaking at or below 80% of system resources.
- Provisioning through LOC-A should complete with no manual intervention.
- The ThinkEdge servers must process RTSP streams with minimal latency, maintaining performance and meeting AI inferencing accuracy benchmarks.

## Validation Results


***Figure 3 – SE350 V2 Performance Results***


***Figure 4 – SE360 V2 Performance Results***


***Figure 5 – SE450 Performance Results***

# Conclusion

The validation successfully demonstrated the efficiency, performance, and interoperability of Lenovo ThinkEdge SE350 V2, SE360 V2, and SE450 servers with Viana, Edge Platform, and LOC-A. Key objectives achieved included:

- Seamless provisioning through Near-Zero Touch Provisioning (nZTP),
- Efficient container orchestration, and
- Robust handling of AI retail inferencing workloads were all met.

All key objectives were achieved, including the onboarding and integration of applications, the setup of application runtime, and managing dependencies across the edge infrastructure.

The tests confirmed stable network performance, optimized resource usage, and reliable connectivity across all components, ensuring smooth deployment and lifecycle management of edge infrastructure.

Furthermore, the ThinkEdge servers processed RTSP streams with minimal latency, maintaining high AI inferencing accuracy under varying loads, meeting performance benchmarks while scaling up. Automated processes for deployment and lifecycle management were efficient, requiring minimal manual intervention.

Overall, the validation demonstrated the capability of the solution to deliver scalable, reliable, and high-performance edge infrastructure deployments, offering clear guidelines for future deployments and performance management.

# Technology Overview

## Hardware and Software Requirements

Viana uses the Edge Platform and Lenovo technologies to enable businesses to run multiple AI use cases on a single device. It integrates various sensors and video feeds for quick deployment of applications like people counting and zone engagement. Viana can track customer behavior and correlate it with POS data across vendors, all managed through a centralized architecture for real-time insights and improved decision-making.

The Edge Platform, integrated with Lenovo Open Cloud Automation (LOC-A), offers edge-native applications, AI models, and tools for building, deploying, and managing edge-AI solutions at scale. The software stack, featuring an OpenVINO AI runtime, optimizes performance across diverse environments. With infrastructure orchestration, zero-touch updates ensure all devices are current through a single interface.

The table below is the system requirements for installing and efficiently running Viana use cases:

***Table 1 – ThinkEdge Systems***

| **Component** | **SE350 V2** | **SE360 V2** | **SE450** |
|---|---|---|---|
| **Processor** | Intel Xeon D-2775TE CPU @ 2.00GHz, 8 Cores, 16 Threads | Intel Xeon D-2775TE CPU @ 2.00GHz, 16 Cores, 32 Threads | Intel Xeon Gold 5318N CPU @ 2.10GHz, 24 Cores, 48 Threads |
| **Memory** | 128GB DDR4 | 256GB DDR4 | 512GB DDR4 |
| **Storage** | 3.84 TB NVMe | 3.92TB NVMe | 3.92TB M.2 & NVMe |
| **OS** | Ubuntu 22.04 LTS | Ubuntu 22.04 LTS | Ubuntu 22.04 LTS |
| **Orchestrator** | Intel Tiber™ Edge Platform 24.11 | Intel Tiber Edge Platform 24.11 | Intel Tiber Edge Platform 24.11 |
| **RTSP Server** | Simulates total number of cameras that can be loaded into the Viana software | Simulates total number of cameras that can be loaded into the Viana software | Simulates total number of cameras that can be loaded into the Viana software |
| **Viana Edge App** | Viana v3.6.1: Performs local inference on RTSP cameras using an RTSP server as dummy cameras. | Viana v3.6.1: Performs local inference on RTSP cameras using an RTSP server as dummy cameras. | Viana v3.6.1: Performs local inference on RTSP cameras using an RTSP server as dummy cameras. |
| **Retail Use Cases** | • **People Counting**<br>• **Anonymous Audience Measurement**<br>• **Traffic Measurement** | | |

*Table 2 - Video Stream Configuration*

| Component | Setting | Notes |
|---|---|---|
| Video Analytic Input Video Stream Parameters | 1920x1080@25fps (1080p)<br>1280x720@25fps (720p)<br>640x480@25fps (480p) | FHD, HD, and SD video streams |
| Number of Input Video Streams for Analytics (Virtual cameras) | 1-n | Each virtual camera stream has high-resolution and low-resolution videos |

# Lenovo Open Cloud Automation (LOC-A)

Lenovo Open Cloud Automation is a framework for automating the configuration and deployment of Lenovo hardware including XCC and UEFI configuration, operating system deployment and inclusion of the node into cloud platforms. LOC-A is designed to do late binding of the hardware to the desired site/cloud to simplify logistics and reduce time to value.

The following diagram depicts the high-level architecture of LOC-A:

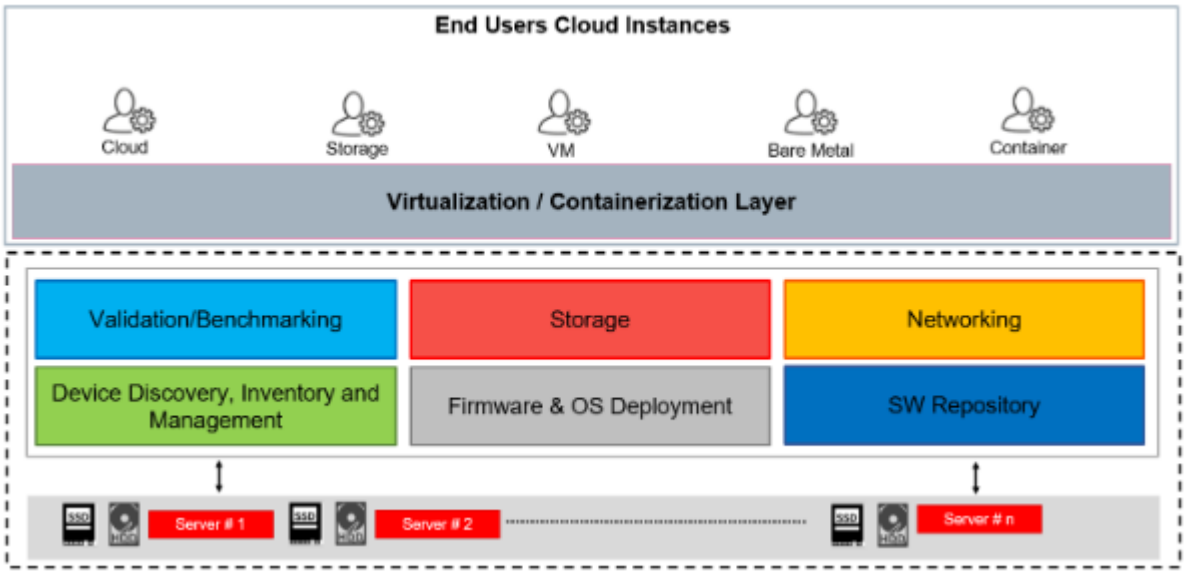


***Figure 6 - The dotted line represents a framework for deploying and managing Cloud and Edge services in a single pane of glass.***

LOC-A supports Kubernetes, Red Hat® OpenShift®, VMware Cloud Foundation™, and bare metal deployments of OSes like Ubuntu, ESXi, and CentOS, simplifying Lenovo hardware deployment at the edge and data center. It integrates partner solutions, like Edge Platform, via plug-ins, extending automated deployment capabilities. LOC-A connects with management platforms through REST APIs, covering VMs, containers, and bare metal. It allows real-time integration of new workloads and scales seamlessly across servers and workloads.

# Intel Tiber Edge Orchestrator

Edge Orchestrator comprises of three sub-components: Edge Infrastructure Manager (EIM), Cluster Orchestrator, and Application Orchestrator. EIM is the fleet management tool for Edge Platform, simplifying lifecycle management for edge devices. It uses a hierarchical data model for policy-based management across organizations, regions, and sites. EIM handles secure, remote onboarding and provisioning with Ubuntu OS but does not interact with Kubernetes Edge Clusters.

Cluster Orchestrator layers on top, managing Kubernetes Edge Clusters and simplifying infrastructure workflows. It provides modular, simple app deployment for developers through its API. It is designed with simplicity and modularity in mind. Application Orchestrator provides the ability to deploy containerized applications across edge node clusters at scale.

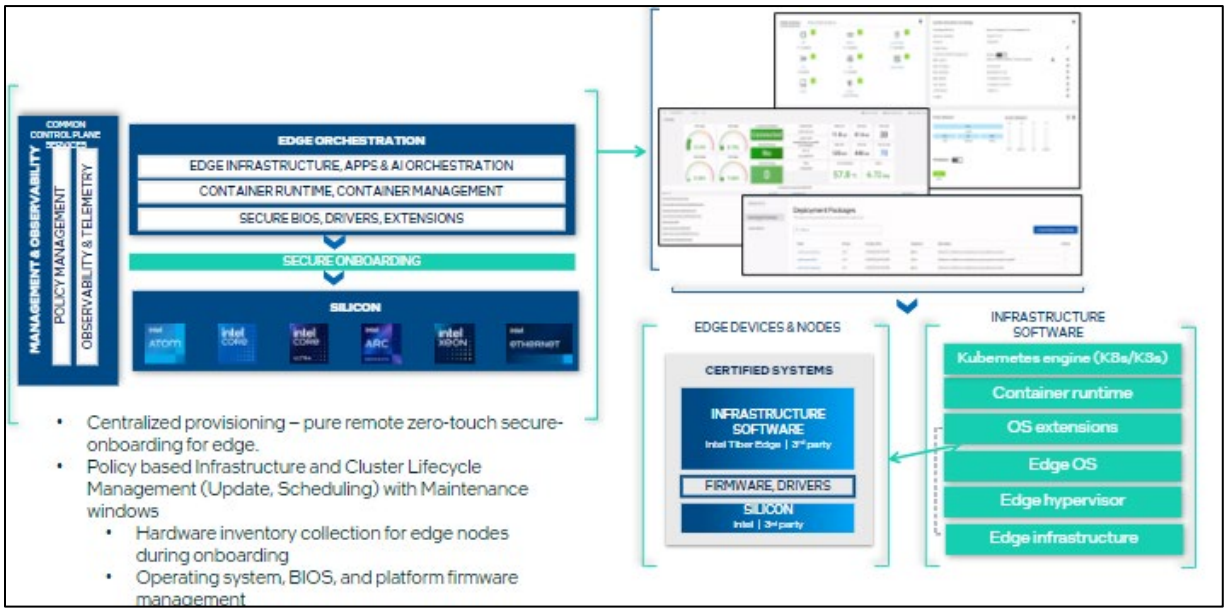


***Figure 7 - Edge Orchestration***

# Viana Advanced Vision Analytics for Data-Driven Insights

Viana is a powerful vision analytics platform that combines IoT hardware and machine learning at the edge. It enables organizations to harness the benefits of machine vision and AI-driven data analysis without the need for dedicated developers.

Viana collects anonymous audience data through high-performance face mapping, ensuring privacy by never storing any personally identifiable information (PII). This allows organizations to gather valuable insights into user behavior while fully respecting privacy regulations.

By analyzing how people work, shop, learn and play in physical and digital spaces, Viana empowers teams to make informed, data-driven decisions that drive better customer experiences and operational efficiencies.

**Key built-in applications of Viana:**

- **People Counting:** Measure foot traffic, including entries, exits, and visit duration, along with demographic data to optimize space utilization and customer flow.
- **Anonymous Audience Measurement:** Capture anonymized visitor data such as age, gender, mood, and time spent within a space. This provides businesses with valuable insights into audience demographics without compromising privacy.
- **Traffic Measurement**: Monitor and analyze vehicle flow in shopping centers or parking districts, capturing details such as make, model, type (e.g., electric vs. fuel-powered), and duration of stay. This data enables retailers to make informed assumptions about customer behavior, distinguishing between longer visits for experiences like cinemas or dining and shorter trips for quick shopping needs. These insights drive strategic decisions on infrastructure improvements, such as adding electric vehicle charging stations or optimizing parking spaces, ensuring the environment aligns with customer preferences and enhances their overall experience.

A common use case for Viana is with customers who already have camera systems in place but seek to unlock deeper insights through advanced analytics. For example, a retailer with an existing camera network can leverage Viana to turn basic surveillance footage into actionable, data-driven insights for smarter decision-making. Viana is compatible with a variety of camera models, from CCTV/POE to USB cameras, providing businesses the flexibility to choose the optimal solution for their specific needs and environment.

While standard cameras capture customer activities, their capabilities are often limited to basic surveillance without providing actionable insights or pattern analysis. Viana, with the power of edge computing, transforms these cameras into intelligent, autonomous systems that not only see but also analyze customer behavior in real-time, enabling data-driven, actionable decisions seamlessly.

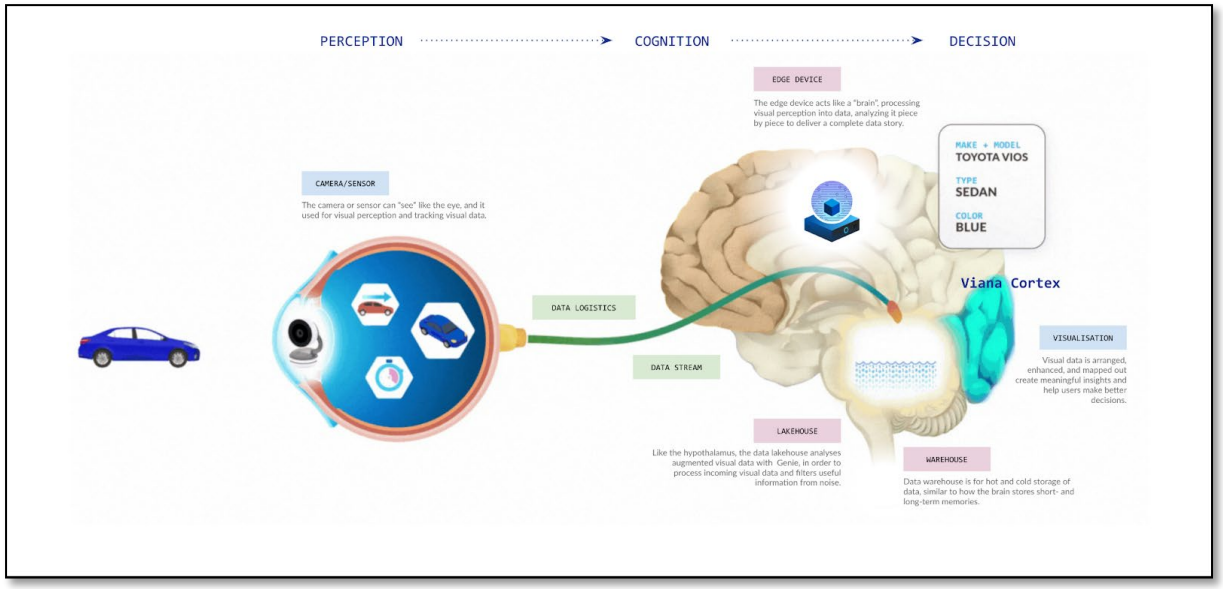


***Figure 8 - Visual Pathway - the "Viana Cortex"***

# Application Design Considerations

## Viana Analytics

Viana prioritizes privacy by ensuring that all analytics are conducted without compromising individual identities. The platform neither captures nor processes personally identifiable information (PII), extracting real-time data from video footage without storing or recording it. AI models are built on synthetic data that simulate virtual 3D environments, not real individuals, and all faces in video inputs are blurred, preventing even support staff from viewing identifiable features. Additionally, a token system anonymizes all data with random numbers, ensuring individuals cannot be traced. Viana performs no facial recognition, keeping all data non-identifiable and fully anonymized, maintaining privacy and security while delivering valuable insights.

## Intel Tiber Edge Platform Security

The Edge Platform delivers Zero Trust security with a comprehensive range of security controls and environment hardening measures, including:

***Table 3 - Below is a summary of security controls and capabilities available from the Edge Platform:***

| Security function | Purpose | Description |
|---|---|---|
| **Zero Trust Architecture** | Identity, Access controls, Policy Enforcement, Device Trust and Defense-in-depth | No user, device, or application is trusted by default, and access is granted based on continuous verification of identity and access controls. |
| **Security Monitoring and logging** | Monitoring, logging, auditing, and alerting | Continuously monitor and log user activities, actions, API calls, and configurations to detect unauthorized access, malware, and security incidents. Generate alerts based on security policies and thresholds to effectively mitigate risks and protect cloud and edge assets. |
| **Security Validation** | Automated Security Testing, Fuzzing and Penetration testing | Continuously evaluate application and system security using automated tools to identify vulnerabilities. Use fuzzing techniques and perform penetration testing to reduce the risk of successful attacks. |
| **Threat Detection and Response** | Vulnerability scanning, Threat identification, Incident response. | Use automated tools to scan source code and binaries for security flaws. Adhere to PSIRTs and InfoSec guidelines to swiftly address breaches and minimize damage. |

*Table 4 - Edge Orchestrator security overview*

| Security function | Purpose | Description |
|---|---|---|
| Identity and Access Management (IAM) | Authentication, Authorization, Role-Based Access Controls, Privileged Access Controls and Zero trust | Protect user, application, deployment, and edge device identities within the Edge Orchestrator. Enforce least privilege through access controls based on entitlements and policies. |
| Data Security | Encrypt data at rest and in transit.<br>Secure storage and Certificate management | Protect data in storage and transit between users, devices, and cloud services from unauthorized access. Securely manage tokens, passwords, certificates, and encryption keys. Use TLS 1.2 & 1.3 for data in transit. |
| Network Security | Firewall,<br>Security Groups,<br>Secure communication,<br>Service Mesh and<br>Certificate manager | Protect the confidentiality, integrity, and availability of cloud and edge data and applications using WAF, TLS 1.2 & 1.3, and cert-manager for certificates. |
| Application Security | Access Controls and<br>Policy Management | Protect the confidentiality, integrity, and availability of Cloud and Edge infrastructure data and applications. |

## System Resiliency

The Edge Platform provides a series of mechanism to ensure continuous operations. The Edge Orchestrator supports high availability. High availability is supported when there are a minimum of three control plane nodes, and the number of control plane nodes is valid (for example, 3 or 5) to avoid split-brain conditions.

## Edge Application, AI Model, and Infrastructure Orchestration

The Edge Orchestrator simplifies edge application deployment and management, making it easier to deploy edge solutions at scale. Edge Orchestrator provides:

- **Secure Infrastructure Management**: Securely onboard and manage edge nodes across sites with zero-trust security for rapid application protection.
- **Orchestration & Automation**: Centrally deploy, update, and configure infrastructure with Edge Orchestrator for automated cluster management.
- **Automated Deployment**: Streamline large-scale remote app installations and updates.
- **Deep Telemetry**: Centralized policy-based lifecycle management and visibility into edge infrastructure.

- **Flexible Configuration**: Define policies and hierarchies to manage physical and runtime environments tailored to your business needs.
- Define policies and hierarchies to manage physical and runtime LOC-A

# Accelerated and Efficient Edge Deployment with LOC-A

LOC-A is a no-code software solution designed to securely expedite the large-scale deployment of distributed edge infrastructure. It enables faster deployment with fewer resources, contributing to reduced CO2 emissions throughout the process. By simplifying and automating complex deployment steps, LOC-A ensures a standardized experience for both infrastructure administrators and field technicians while remaining adaptable to specific customer, deployment, and site needs.

**Primary Functions of LOC-A:** LOC-A streamlines the near-zero touch provisioning and onboarding of edge nodes by enabling late binding of edge servers to their owners. It automates the provisioning and configuration of Day-0 settings before onboarding, allowing the edge nodes to host cluster software and applications seamlessly, with the base OS deployed remotely.

**Infrastructure Automated Deployment**: LOC-A automates the installation and configuration of distributed edge infrastructure, from bare metal to production-ready IaaS/CaaS, ensuring rapid and secure deployment by applying consistent configuration patterns across hardware, OS, and clusters.

**Scalability**: It efficiently scales edge deployments, allowing the addition or removal of nodes as needed, from small to large-scale operations, leveraging Kubernetes for flexibility.

**Lifecycle Management**: LOC-A provides comprehensive lifecycle management for edge infrastructure, including deployment, redeployment, and cluster expansion, ensuring consistent and reliable edge service operations.

**Orchestration and Integration**: The platform integrates with edge-cloud management and orchestration tools, enabling automated workflows and full-stack edge application enablement for enhanced user experience.

By enabling rapid, large-scale provisioning without golden images or staging environments, LOC-A minimizes transport emissions and supports IT decarbonization.

## Performance Requirements

- **Network Connectivity**: The software should be capable of handling high volumes of data transfer across the network without bottlenecks. It must efficiently manage data transmission rates to support seamless communication and data exchange, ensuring that the network infrastructure can sustain peak loads and maintain performance levels. It is recommended to have at least 50mbps speed upload and download bandwidth. Also, the system must be capable of handling disconnection of edge devices and recover automatically.

- **Real-time Processing:** The software must ensure rapid response times to user actions and system requests. The goal is to minimize latency, providing near-instantaneous feedback to enhance user satisfaction and efficiency. Typical response time targets should be defined based on specific use cases and industry standards.

## Scalability Requirements

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

# Architectural overview

## Introduction

The Edge Platform integrates with LOC-A and Lenovo XClarity Administrator comprising of cloud native functions and services, that allow organizations to build, run, secure, manage, connect, and protect edge-native applications at both the near and far edge locations.

The Edge Platform architecture consists of both virtualized infrastructure and applications to deliver a comprehensive managed edge stack with Lenovo and Viana. Underlying technologies from key cloud native applications help reduce the complexity and increase the support capabilities through re-use of industry standard products.

Edge Platform framework facilitates the large-scale deployment of Viana services which performs the inferencing on the video feeds from the cameras across multiple edge nodes. The life cycle management of the applications is performed by the Edge Orchestrator while the platform and hardware configuration are performed by the LOC-A Orchestrator.The overall solution shown in Fig. 4 illustrates both hardware and software elements involved in the end-to-end data processing and inferencing on the camera feeds. The main aspects of the solution include:

### Edge Node Provisioning:

- LOC-A and Edge Platform work together to provision the edge node platform. LOC-A handles the BMC/UEFI configuration as well as installation and configuration of the OS on edge nodes, preparing them for seamless integration with the Edge Orchestrator.
- After the edge node OS is provisioned, Edge Orchestrator synchronizes its Edge Infrastructure Manager (IM) database with the LOC-A database, registering the provisioned edge nodes and their OS instances. It then provisions the necessary services and application runtime on the edge nodes.

### Application Provisioning:

- Edge Orchestrator hosts the application images, such as Viana containers, in a registry, that could be deployed at the edge nodes once the OS provisioning is performed by LOC-A.
- Edge Orchestrator also performs the resource allocation for the application at each node, as well as assumes the role of the complete life cycle management of the applications running on the edge node such as load balancing, migration, and backup services are performed.
- The orchestrator also configures the telemetry of the edge node state to perform the real-time control of the edge node operations to meet the quality of service and service level agreements as needed for the end-to-end applications.

### AI Inference Pipeline:

- The data originated at the camera is configured to be routed to the applications running at the edge node. The AI inferencing application, Viana uses OpenVINO for accelerating the inference-based compute on the Intel processors of the edge node.

- The AI-Model is typically generated at the cloud, and in this instance, using the Azure based Viana cloud application, and then applied to each node for the inferencing pipeline.
- The results of inferencing are displayed in the dashboard of the Viana application.

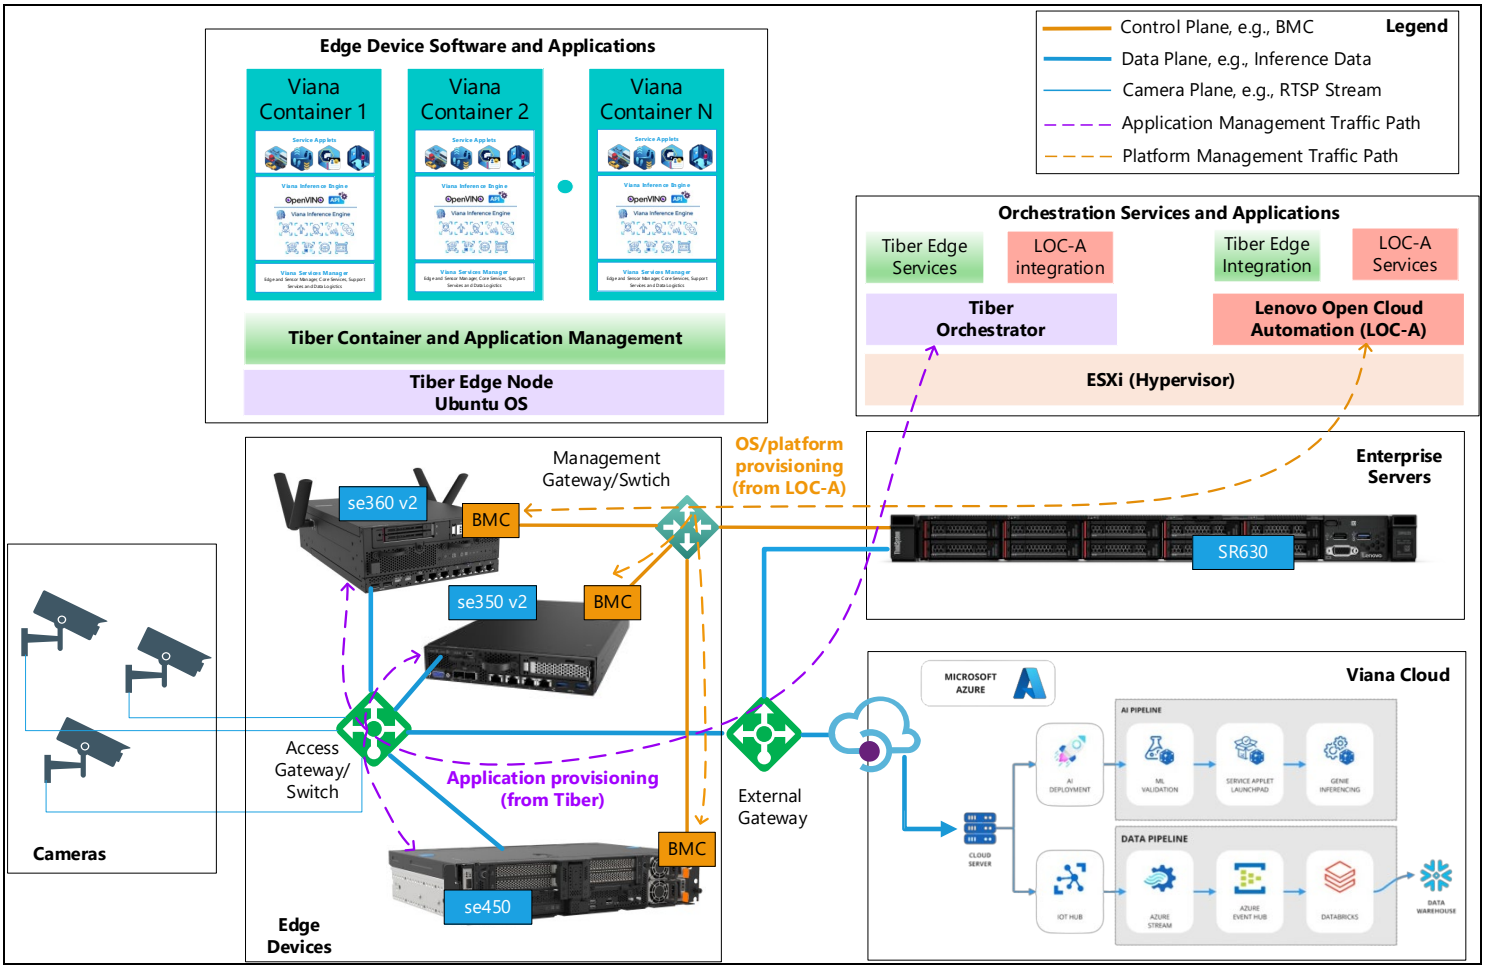


**Figure 9 - Overall Solution Design: Edge Orchestrator streamlines the provisioning and deployment of Viana applications.**

# Component Model

## Intel Tiber Edge Orchestrator

The Edge Orchestrator, a key component of the Edge Platform, manages the administrative functions of setting up the infrastructure, packaging and deploying applications, Identity management and role-based security and monitoring of all elements. The Edge Orchestrator utilizes a collection of both Intel developed and Industry Standard cloud applications to manage:

- Edge Node clusters
- Application catalog(s) and Deployments
- Platform identity and management services
- Observability and reporting

The Edge Orchestrator streamlines the deployment and administration of edge software and hardware by managing the various aspects of both the HW and software lifecycles.

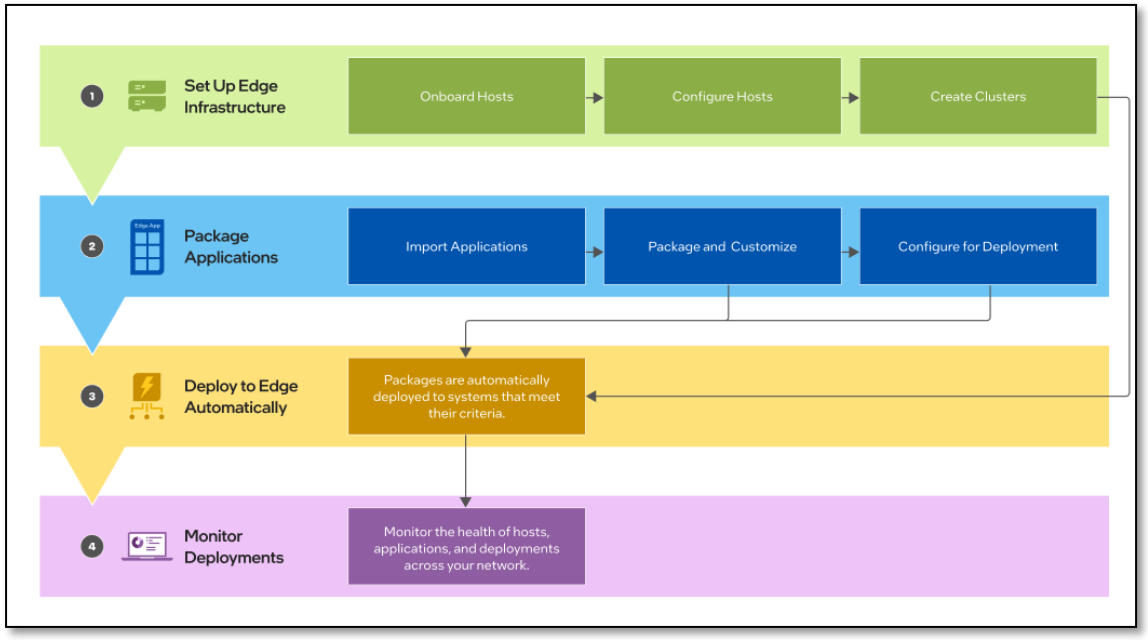


***Figure 10* – Intel Tiber Edge Orchestrator - Pipeline to Production Workflow**

### Edge Node

The Edge Node is a secure server that hosts applications at the edge, orchestrated using the Edge Orchestrator. The Edge Node consists of industry standard containerized applications and microservices that optimizes business results at the edge.

### Edge Node Cluster(s)

The Edge Node onboarded to the Edge Orchestrator is designed to run as a cluster, which helps improve performance and availability at the customer location. However, not all locations require extra capacity or high availability, so Intel has developed the concept of two types of Edge Nodes:

- Single Node Cluster – 1 server that delivers all the applications required for the location.

- Multi-Node Cluster – 2 or more servers that are part of the same control plane and the Edge Orchestrator can automatically deploy application to the right host depending on the application needs.

## Orchestrating a Retail Edge Infrastructure using Intel Tiber Edge Platform

The Edge Orchestrator plays a key role in managing and orchestrating distributed edge and AI applications across store locations. It oversees AI models, applications, and container infrastructure, simplifying orchestration through helm charts and managing application lifecycles, including upgrades and telemetry collection. In this use case, retail store edge applications that enhance customer experiences and streamline operations are orchestrated using the Edge Orchestrator, across distributed edge locations.

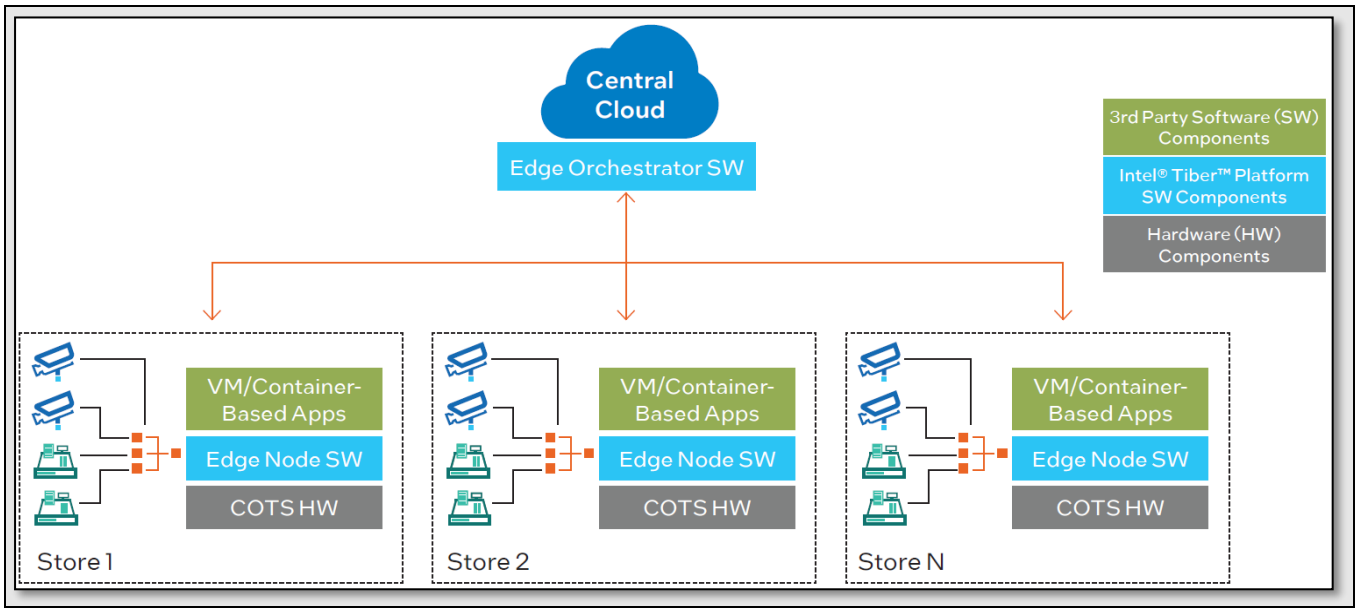


***Figure 11 -*** Distributed Retail Stores - Edge AI Orchestration

Deployment metadata related to edge locations help with creating policies for automatic deployment of such workloads on clusters on new locations come online and connect to the Edge Orchestrator. The containerized application endpoint becomes accessible through the Edge Orchestrator.

## Edge Orchestrator Dashboard

Edge Orchestrator dashboard gives a high-level, single-pane-of-glass status of all the hosts, and application deployments.

For example, in the following host status, there are a total of 3 hosts. All the three hosts are running.

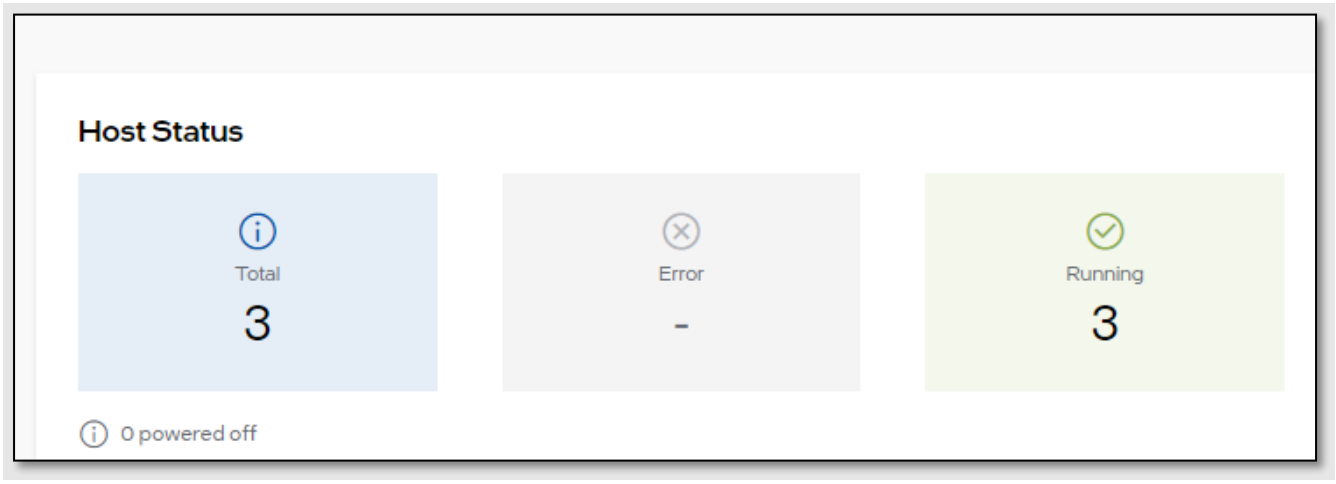


***Figure – 12 View of Host Status***

# Lenovo Open Cloud Automation Component Model

## Near-Zero Touch Provisioning of Edge Nodes

Lenovo Open Cloud Automation's (LOC-A) 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.

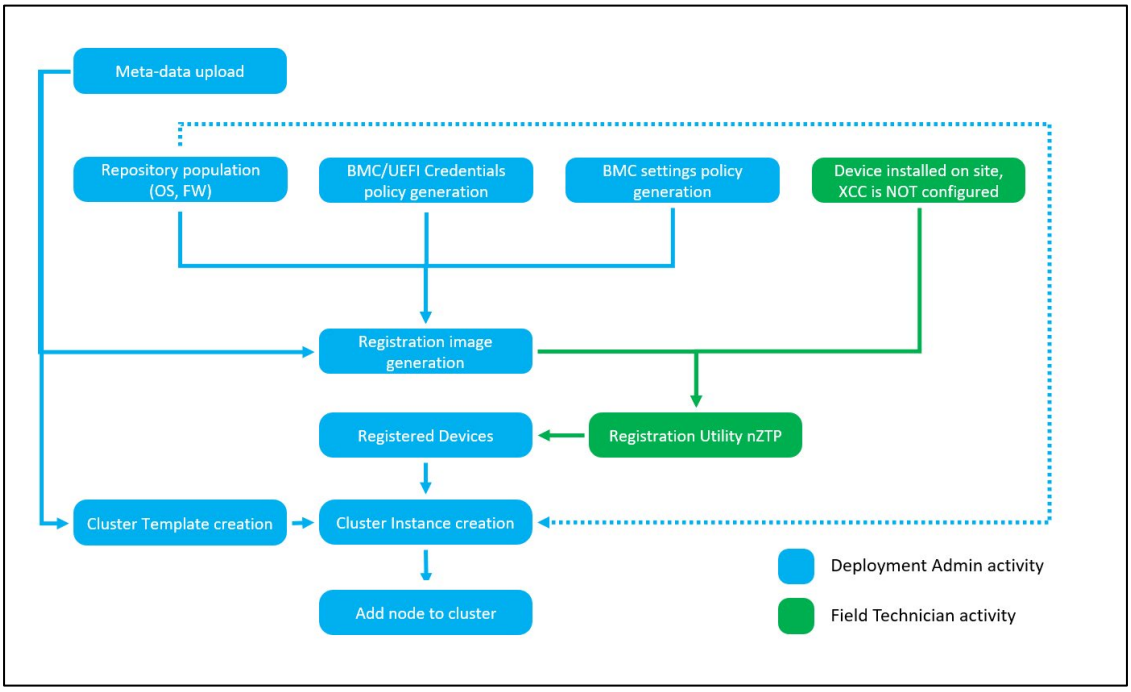


***Figure 13 – LOC-A workflow model for nZTP and OS deployment.***

## 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 flavors. 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.

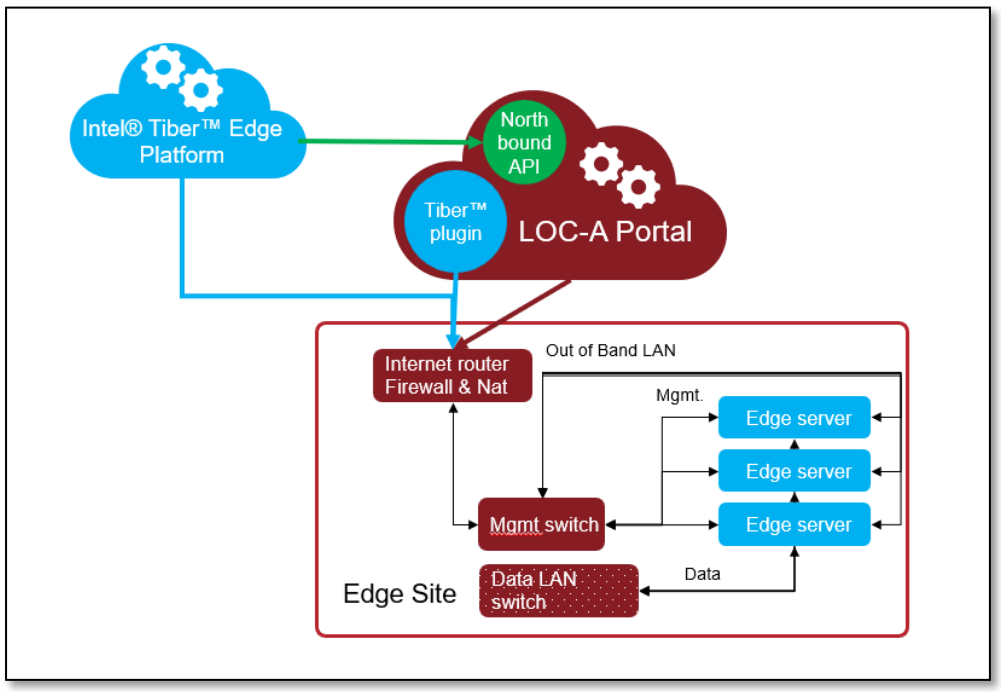


***Figure 14 – LOC-A Plugin mechanism and northbound API***

# Viana Component Model

## Connecting Cameras to Edge AI Computing

- **Evaluate Current Infrastructure**: Thoroughly review the existing camera setup, including camera types, network architecture, and power supply, to identify any compatibility or performance gaps.
- **Determine Site Requirements**: Collaborate with partners to establish optimal camera placement and field of view (FOV) for accurate data capture.
- **Whitelist Viana APIs and Endpoints**: Once the camera model is confirmed, whitelist necessary ports and services to enable the Viana platform.
- **Select Suitable Edge Devices**: Choose edge devices compatible with the current infrastructure that can process real-time video feeds to meet system requirements.
- **Install Edge Devices**: Place edge devices (e.g., Edge Platform) near cameras or centrally, as per the architecture, to enhance performance and coverage.
- **Configure on Viana Portal**: Set up edge devices and cameras in the Viana portal for efficient video feed processing and event detection using AI and machine learning.
- **Integrate with Existing Systems**: Seamlessly connect edge devices with current systems, such as video management software (VMS) or other sensors.

## Viana: Maximizing Edge AI Computing

Viana maximizes the potential of edge AI computing through visual analytics by offering:

- **Advanced AI & ML Capabilities**: Equips edge devices with sophisticated visual AI and machine learning to detect and analyze events accurately. It can stack multiple AI models, enabling seamless communication between modules and centralized data visualization, allowing businesses to monitor diverse metrics on a single dashboard.
- **Scalability & Flexibility**: Provides a scalable framework that easily adapts as edge devices are added or removed.
- **Built-in Privacy Tools**: Ensures privacy by not storing or recording any personally identifiable information (PII), simplifying data security for users.

Viana empowers organizations with actionable insights for informed, data-driven decisions, enhancing operational efficiency, profitability, and customer experiences.

## Viana Deployment Architecture Modes

Viana offers flexible deployment options across multiple architecture modes, including full edge-AI, hybrid (a combination of Edge-AI and Cloud-AI), and a “gateway” model (edge connector with Cloud-AI). Each of these deployment modes aligns with different Viana customer packages, providing tailored solutions to meet various operational needs. Refer to the figure below for a detailed breakdown of these architecture modes.

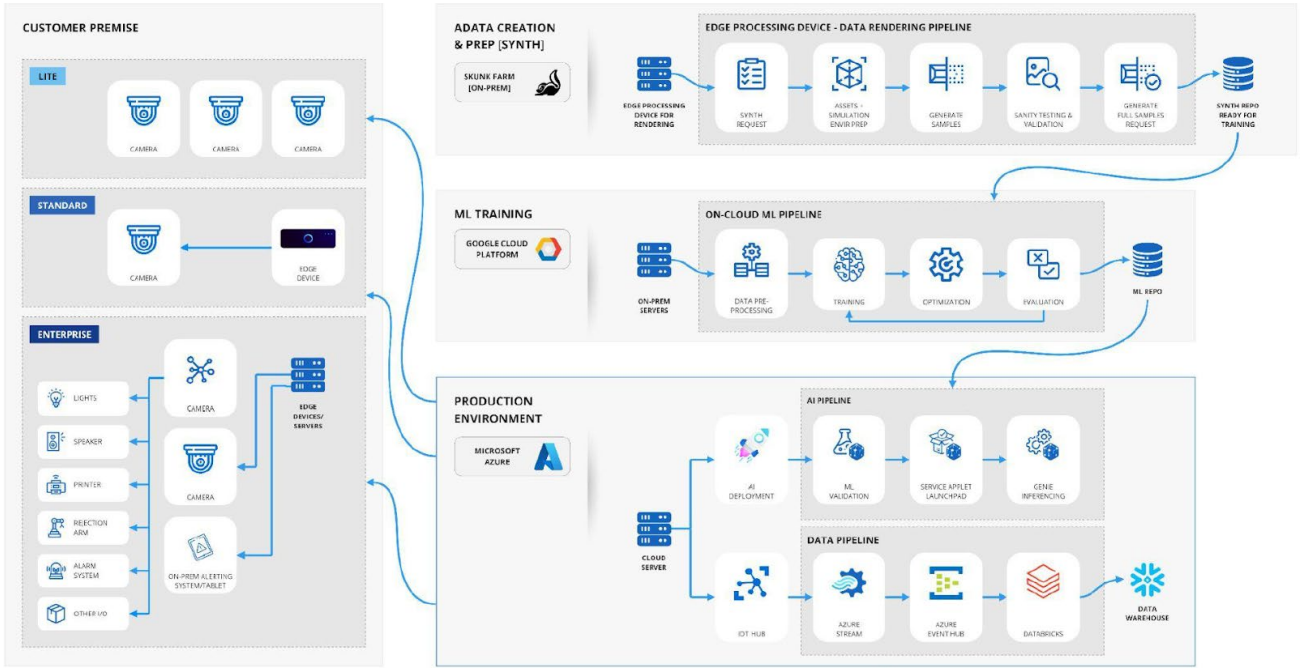


**Figure 15 - Viana Solution Architecture**

# Deployment Model

## Viana License Provisioning and access to the dashboard

After purchasing licenses, the primary business contacts will receive an email invitation to set up their account. Follow the instructions provided in the email to complete the setup. Viana is primarily accessed through its web-based portal at https://portal.viana.ai.
Ensure that both the edge device and deployed cameras are online to start tracking visual data.
For more information on the portal and its features please visit our Helpdesk article.

## Edge Platform Host Observability Dashboard

Leveraging Grafana, the Host Observability Dashboard, from the Edge Platform provides a view of the general performance, system information, and track the connectivity and maintenance status of the host.

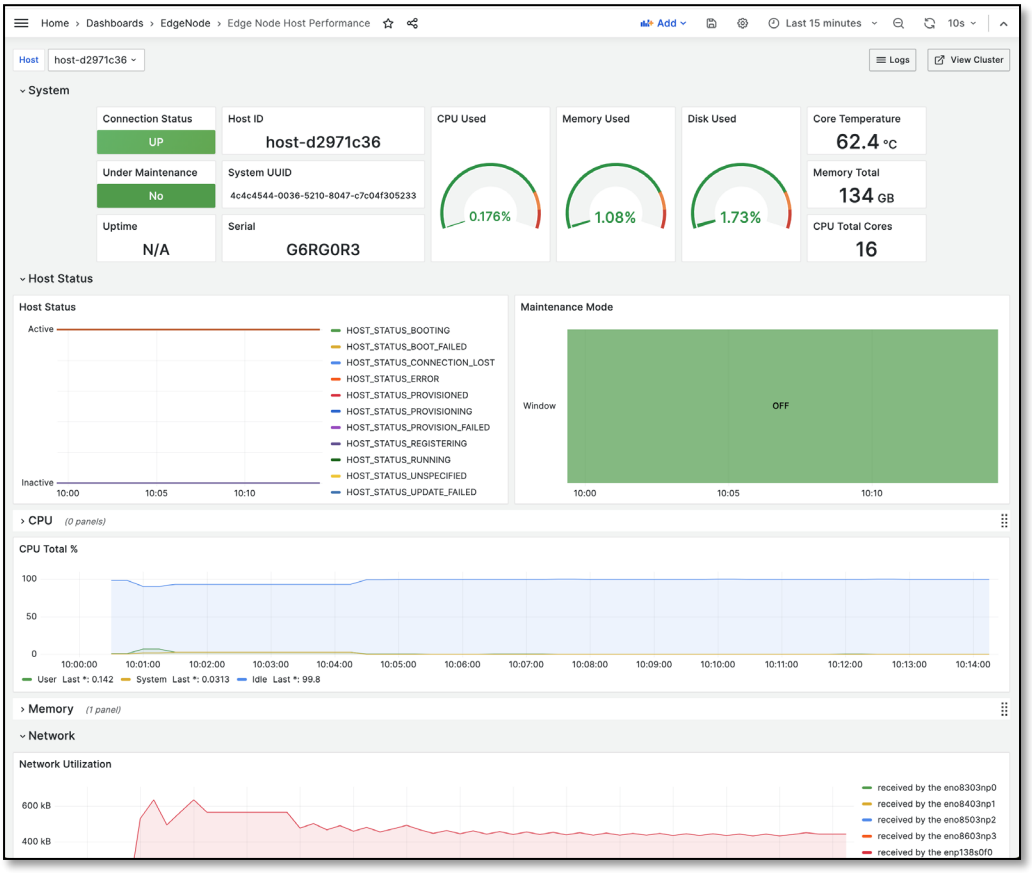


***Figure 16 - Edge Platform Host Observability Dashboard***

### View Host Logs

From the Host Performance dashboard, click the "Logs" button in the top right to access a dropdown of available logs for the Edge Node.

### Cluster Observability

If a cluster is deployed on the Edge Node, click "View Cluster" in the top right of the Host Performance dashboard to access cluster telemetry, including logs and resource usage for both the cluster and any applications running on the Edge Node.

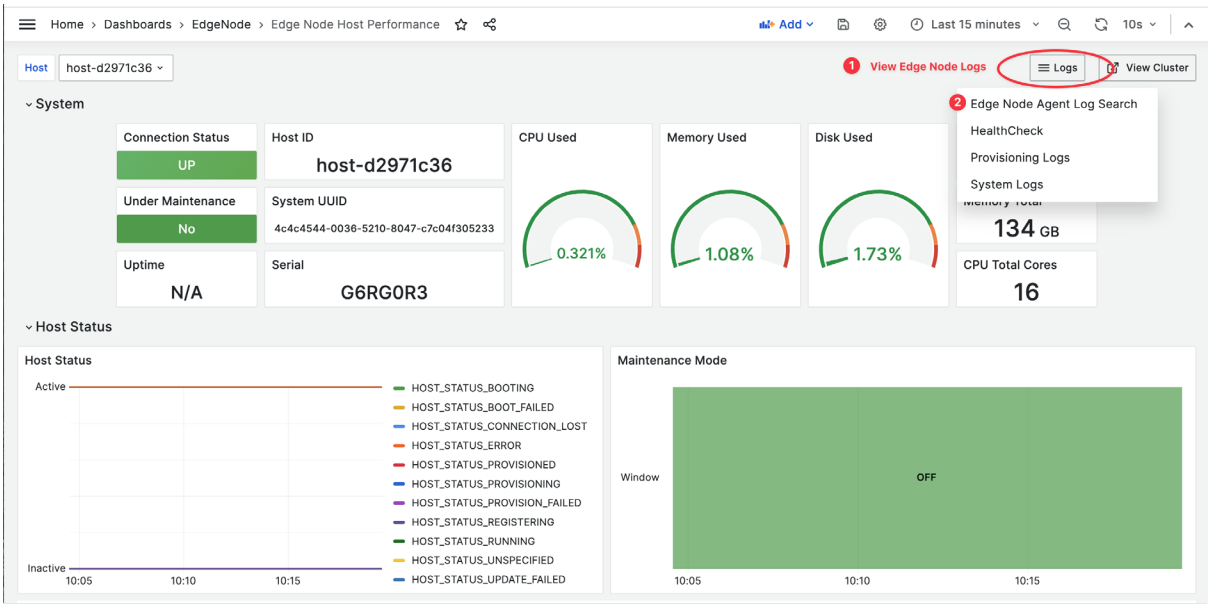


***Figure 17 - Cluster Observability***

## Lenovo Open Cloud Automation Portal

LOC-A simplifies the deployment of edge nodes at scale by enabling users to create deployment templates which can be used to deploy multiple edge nodes, at multiple sites, in parallel. In this section is a high-level overview of the steps required to create an instance via the LOC-A Portal. For full details, please consult the Lenovo Open Cloud Automation User Guide.

## Populating Site Metadata and OS Images

Preparing site metadata in LOC-A involves creating an Excel spreadsheet that includes details for Sites, IP Ranges, Network Services, and Cloud Services used by the edge sites. Lenovo provides a template of the spreadsheet with documentation, examples for guidance, and built-in validation features to simplify configuration.

If a Lenovo XClarity Administrator (LXCA) system is registered as a cloud service for a site, all systems registered in LOC-A for that site are automatically managed by LXCA, eliminating the need to on-board nodes into LXCA manually. LXCA also catalogs all telemetry data from managed nodes.

Once completed, the spreadsheet is uploaded to the LOC-A Portal to populate data into the DCIM database. After a successful upload, the site(s) and associated resources will appear in the ‘Setup’ section tabs of the LOC-A Portal.

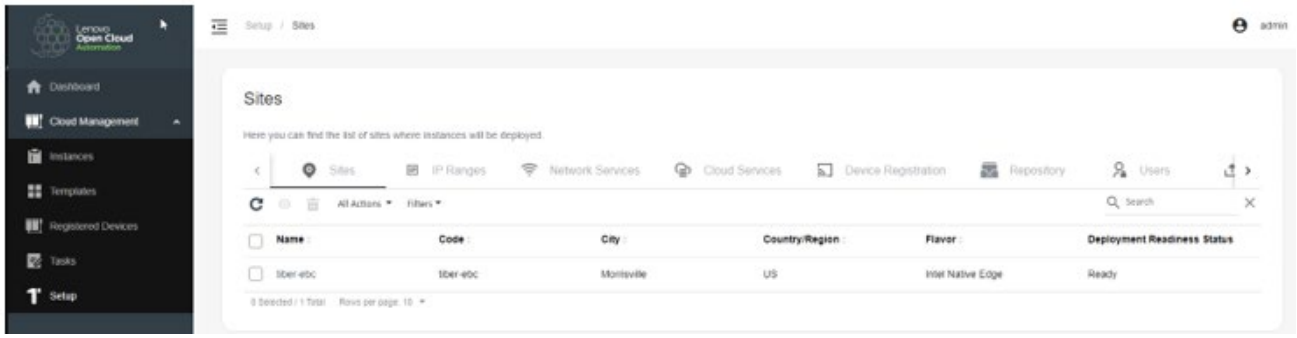


***Figure 18 – LOC-A Portal with metadata populated***

To deploy edge nodes, upload an OS image to the LOC-A Portal's image repository. Once the Edge Platform team's ISO is uploaded and validated, it will appear under the 'Repository' tab in the Setup section.

## LOC-A Templates

LOC-A deploys multiple sites and edge nodes in parallel using templates. Each template specifies the deployment flavor (e.g., "Intel Native Edge" for Edge Platform), networking details, device type, minimum nodes, and authentication configuration.

## Registering Nodes with LOC-A

Hardware can be registered in LOC-A via discovery, manual IP entry, or near Zero Touch Provisioning (nZTP). This Lenovo Validated Design prioritizes nZTP for optimal edge node deployment. An edge technician receives a registration image from the LOC-A Portal, installs the nodes, and completes registration using the LOC-A Registration utility.

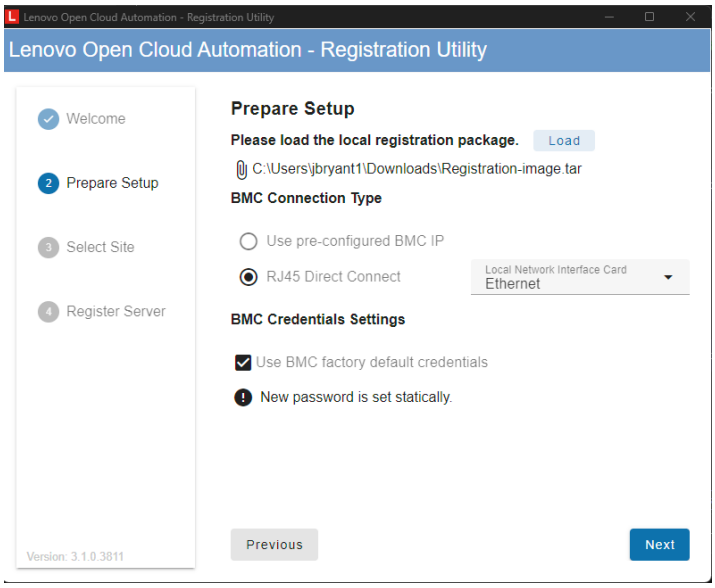


***Figure 19 - Registering Nodes with LOC-A***

The LOC-A Registration Utility allows an edge technician to connect their laptop to the RJ45 XCC management port of the target edge node. After loading the registration package and selecting the network, the technician chooses the site and starts the registration process. The XCC is configured with network settings and credentials, then the edge node is booted, and network services are validated. Upon successful completion of the checks, the node is registered with LOC-A and the status is confirmed on the 'Registered Devices' page in the portal.

## Instance Creation

When a LOC-A template is used to deploy a node or create a cluster, it generates an instance of that template. Instances are created from the 'Instances' page using templates from LOC-A's inventory. LOC-A will then evaluate the defined sites to determine which ones meet the deployment criteria for the selected template. The datacenter administrator can then choose one or more of the eligible sites for creation of the instance(s).

# Operational Considerations

## Systems management

The systems management approach for this solution integrates advanced lifecycle management, software distribution, firmware updates, and application management processes to ensure high availability and optimal performance. Leveraging Lenovo Open Cloud Automation (LOC-A), Edge Platform, and Viana, the solution maintains target service levels by automating key management functions across the edge infrastructure.

## Lifecycle Management

The **Lenovo Open Cloud Automation (LOC-A)** platform plays a central role in provisioning, configuring, and managing the lifecycle of the Lenovo ThinkEdge SE350 V2, SE360 V2, and SE450 devices. LOC-A ensures seamless deployment from initial provisioning through to decommissioning by automating processes such as:

- **Near Zero-touch provisioning (nZTP)**: Devices are automatically onboarded, reducing the time and manual effort involved in configuring systems.
- **Firmware management**: LOC-A enables firmware update at instance creation time, ensuring that all edge devices start with the latest firmware releases across the environment without manual intervention.
- **Health monitoring, alerting and firmware updates**: When Lenovo XClarity Administrator is registered as a cloud service, registered nodes are automatically on-boarded to ensure constant monitoring of system health with proactive alerts for potential issues, including firmware updates, reducing downtime.

### Software and Firmware Updates

To maintain optimal performance and security, the system utilizes both **LXCA** and **Edge Platform** for handling software and firmware updates. The update process is designed to be seamless and minimally disruptive:

- **Firmware updates via LXCA**: Devices are updated with the latest firmware versions directly through LXCA, ensuring consistency across the infrastructure. Updates are scheduled during maintenance windows to avoid service interruptions.
- **Software updates and patches**: Edge Platform provides centralized software management tools, so that new versions of system-level software can be distributed efficiently across all nodes and scheduled. This includes updates for both operating systems and containerized environments.

### Application Management and Updates

The collaboration leverages **Edge Platform** and **Viana** to manage application updates and deployment processes, focusing on real-time AI inferencing and data analysis workloads.

- **Container orchestration with Intel Tiber Edge Platform**: The Edge Platform manages the deployment and scaling of AI workloads running in containers on the ThinkEdge devices. The orchestration ensures that application updates are deployed with zero downtime, using rolling updates to apply changes to containers without impacting services.
- **Viana application updates**: Viana provides regular updates to its AI inferencing models and data analytics capabilities. These updates are pushed through the containerized environment managed by the Edge Platform. This ensures that applications remain up-to-date with the latest AI capabilities without requiring any manual intervention.
- **Automated rollback and version control**: Both Intel's Edge Platform and Viana integrate automated rollback processes in case an update leads to service degradation, ensuring continuity of business operations.

### Backup and Recovery

In case of system failures, the solution incorporates a comprehensive backup and recovery strategy:

- **Data backup**: All critical data, including configuration settings and AI inference results, are regularly backed up to a secure storage location.
- **Configuration management**: LOC-A handles configuration backup, ensuring that system settings can be restored quickly in case of hardware failure or system misconfiguration.
- **Disaster recovery**: In the event of a catastrophic failure, the backup system supports a full recovery, minimizing downtime and ensuring that systems can return to operational status rapidly.

### Change Control and Configuration Management

Managing changes across the system, including hardware, software, and applications, is handled by:

- **Lenovo XClarity Administrator**: LOC-A works in concert with XClarity to enable configuration management. It allows admins to track changes across the infrastructure, including ThinkEdge devices, and implement change control policies.
- **Edge Orchestration**: Any changes to containerized applications are managed centrally through the Edge Orchestrator dashboard from Edge Platform, ensuring consistent configuration across all nodes and enabling the deployment of updates or patches without configuration drift.

## Server / Compute Nodes

The solution utilizes **Lenovo ThinkEdge SE350 V2**, **SE360 V2**, and **SE450** servers as the primary compute nodes. These edge-optimized servers provide the necessary performance and reliability for AI inferencing and data analysis at the edge. All nodes in this configuration are homogeneous, offering a unified hardware environment to simplify provisioning, management, and updates via LOC-A/LXCA and Edge Platform.

- **Lenovo ThinkEdge SE350 V2** and **Lenovo ThinkEdge SE360 V2**: Designed for compact edge deployments, these servers provide compute power with scalability for AI workloads, while maintaining energy efficiency and operational resilience in diverse environments.

- **Lenovo ThinkEdge SE450**: Acts as a high-performance node, particularly for handling larger AI inferencing loads and simulating RTSP streams for IP camera workloads in this collaboration.

This uniformity across the compute nodes ensures seamless integration with LOC-A for lifecycle management and Edge Platform for containerized application orchestration.

# Networking

## Intel Tiber Edge Platform

Edge Platform Connectivity Requirements:

- **Network Connectivity**: Edge Platform requires direct connectivity with LOC-A and in the general connection should be able to handle reasonable amount of traffic to allow a smooth user experience. No special requirements for the edge-node software which should be able to reach through the OS network the Edge Orchestrator
- **Edge Platform Services Whitelisting**: Whitelist relevant Edge Orchestrator APIs on the end points that are mentioned in the documentation. This includes additional configuration such as DNS and firewall appliances.

## LOC-A Platform

A typical network configuration for Lenovo Open Cloud Automation is show below:

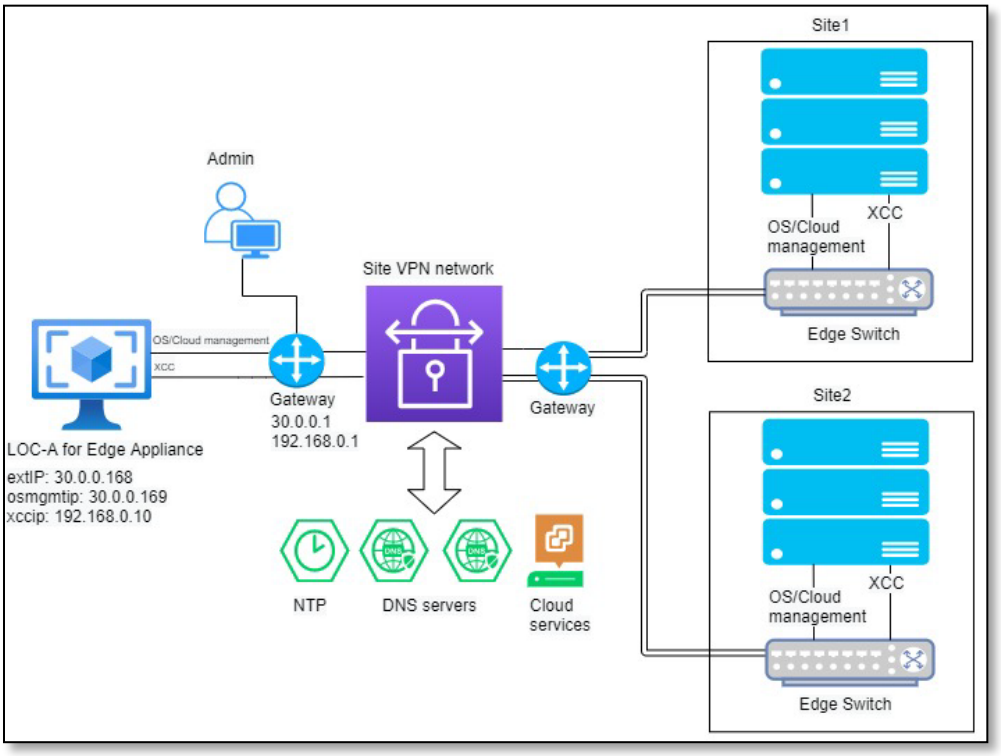


***Figure 20 – Network topology***

LOC-A supports either a dedicated edge Out of Band (OOB) network separated from cloud networks, or a

Layer 3 network on which OOB and cloud networks can be shared. For Edge Platform deployments using LOC-A for edge node deployment, VPN connectivity from the LOC-A Appliance to the target edge nodes is required. In addition, two IP ranges are required to be defined to support automatic configuration of networking for the XCC as well as the edge node's host networking connection.

There are a few ports that need to be whitelisted between the LOC-A Portal and the XCC and OS/Cloud Management networks. The table below indicates these ports and their purpose (note that this assumes that other protocols like DNS and NTP are already whitelisted).

***Table 5 – Port and Purpose***

| Port Number | Purpose |
|---|---|
| 443 | Secure communication between the LOC-A Portal and Registration Utility |
| 13001 | Secure connection for Confluent (OS Deploy) communication |
| 1900 | System discovery |

If doing discovery of nodes, It is strongly recommended that IPv6 be enabled on the XCC network and that the XCC have IPv6 enabled. Confluent's full discovery capabilities will not function without it.

## Viana

- **Network Connectivity**: Connection should be capable of handling high volumes of data transfer across the network without bottlenecks. It must efficiently manage data transmission rates to support seamless communication and data exchange, ensuring that the network infrastructure can sustain peak loads and maintain performance levels. It is recommended to have at least 50mbps speed upload and download bandwidth. Also, the system must be capable of handling disconnection of edge devices and recover automatically.
- **API/Viana Services Whitelisting**: Whitelist relevant Viana APIs on the end points, this will be provided by the Viana team according to the different modules purchased.

# Summary

This solution incorporates a robust and automated approach to systems management, lifecycle management, and application updates to ensure high performance and availability across Lenovo ThinkEdge devices and meldCX applications. By integrating, Edge Platform, LOC-A, XCC, and Viana, the solution streamlines processes like provisioning, firmware updates, and container orchestration, while also enabling seamless deployment of AI workloads. With proactive monitoring, automated updates, and backup and recovery mechanisms, the system is designed to minimize downtime and optimize resource utilization. Network connectivity is essential, requiring specific configurations for seamless communication between LOC-A, Viana, and Intel Tiber Edge Platform. This comprehensive infrastructure ensures scalability, efficiency, and resilience in handling evolving business needs and workloads.

# Appendix: Lenovo Bill of materials

***Lenovo ThinkEdge SE350 V2 Bill of Materials***

| Description | Part Number | Qty |
|---|---|---|
| ThinkSystem M.2 N-30m2 480GB Read Intensive NVMe PCIe 3.0 x4 NHS SSD (Industrial) | 03LC295 | 2 |
| ThinkSystem 7mm U.3 7450 PRO 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD SED | 03KH365 | 1 |
| System Board, ICE Lake D-2775TE , with TPM, with battery - Taiwan | 03LC721 | 1 |
| SE350V2 Node Label Group | 03LF199 | 1 |
| SE350V2 intrusion holder | 03LE725 | 1 |
| SE350V2 CPU Air Baffle | 03LE425 | 1 |
| SE350V2 7mm Cage Instruction Label_Lenovo International | 03LF194 | 1 |
| SE350 V2 Top Cover Bracket | 03LF183 | 1 |
| SE350 V2 Node Chassis(Without Top Cover) | 03LF189 | 1 |
| SE350 V2 Front Input/Output bezel(25G/10G) | 03LF187 | 1 |
| Power Supply Liteon 500W 230Vac/115Vac AC Non-Hot Swap | 03LE826 | 1 |
| Power Input Board, L Shape, Rear, 1U AC 110-220V PIB | 03LF052 | 1 |
| NODE BOT RUBBER PAD | 03LE726 | 1 |
| MECHANICAL,SE350V2 Cable wall | 03LD386 | 1 |
| MECHANICAL,SE350 V2 7mm HDD Cage Assembly | 03LD383 | 1 |
| Label,SE350V2 node Self Service Label_Language International | 03LC858 | 1 |
| Interposer, 1U I/O Board Interposer Card, Internal, 1U I/O Board to System Board, 1U Bridge Board | 03KM995 | 1 |
| I/O FILLER Group Bill of Material (GBM) | 03LE692 | 1 |
| Front I/O Assembly, Front, for SE350 V2 with LEDs (Power, ID, Status, Activation) | 03KM994 | 1 |
| Front I/O assembly, Front , For SE350/360 V2 with 4 x 25Gb SFP28 /10Gb SFP and 2 x 2.5Gb RJ45, 1U I/O Board | 03KM992 | 1 |
| For SE350/ SE360 V2 USB-C to VGA Adapter | 03LF889 | 1 |
| FAN, 40mm 25K RPM, non-hot swap, single rotor | 03KM953 | 3 |
| Cable,internal,150mm,SE350 V2, MB to OP Panel, 2x10-1x15 | 03LF066 | 1 |
| Cable, internal, 85mm, SE350 V2, IO board to HDD BP1, power, 2x6-2x6 | 03LF085 | 1 |
| Cable, internal, 300mm,SE350 V2, MB to HDD BP2 Gen4, MCIO 74-74 | 03LD285 | 1 |
| Cable, internal, 250mm,SE350 V2, MB to HDD BP1 Gen4, MCIO 74-74 | 03LD284 | 1 |
| Cable, internal, 100mm, SE350 V2, MB to IO board, power,2x2-2x2 | 03LF079 | 1 |
| Cable, internal, 100mm, SE350 V2, MB to intrusion switch, 1x3-1x2 | 03LE418 | 1 |
| Backplane, SATA/NVMe, 2.5", 2 Bay, Front, Any Bay, 1U 7mm Backplane 2 | 03KM997 | 1 |
| Backplane, SATA/NVMe, 2.5", 2 Bay, Front, Any Bay, 1U 15mm/7mm Backplane 1 | 03KM996 | 1 |

| | | |
|---|---|---|
| 64GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM | 02JK971 | 4 |
| 1.5m, 10A/100-250V, C13 to C14 Jumper Cord | 39M5375 | 3 |

***Lenovo ThinkEdge SE360 V2 Bill of Materials***

| Description | Part Number | Qty |
|---|---|---|
| ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD SED (with Heatsink) | 03KH324 | 2 |
| ThinkSystem M.2 7450 PRO 1.92TB Read Intensive NVMe PCIe 4.0 x4 NHS SSD SED (with Heatsink) | 03KH325 | 2 |
| System Board, ICE Lake D-2775TE , with TPM, with battery - Taiwan | 03LC721 | 1 |
| SE360V2 Top Cover Assembly | 03LF294 | 1 |
| SE360V2 PMB Air Baffle AC | 03LE427 | 1 |
| SE360V2 PCIE dummy filler Low profile | 03LF414 | 1 |
| SE360V2 intrusion holder | 03LE733 | 2 |
| SE360V2 FAN SWITCH kit | 03LE731 | 1 |
| SE360V2 Bottom AirBaffle | 03LE429 | 1 |
| SE360 V2 Node Bottom system service label_Lenovo International | 03LF408 | 1 |
| SE350 M.2 Heat Sink | 02JJ078 | 1 |
| Riser Card, Gold Finger, 2U, M.2 GenZ 4C Connector, Riser 1 | 03KM983 | 1 |
| Power Supply Liteon 500W 230Vac/115Vac AC Non-Hot Swap | 03LE826 | 1 |
| Power Input Board, L Shape, Rear, 2U AC 110-220V PIB | 03LF050 | 1 |
| Non-RAID, PCIe Gen 4, Quadruple M.2 NVMe, 2U M.2 Adapter | 03KM984 | 1 |
| Non-RAID, PCIe Gen 3, Quadruple M.2 NVMe, 2U M.2 Cabled Adapter | 03KM985 | 1 |
| NODE BOT RUBBER PAD | 03LE726 | 1 |
| Miscellaneous Parts Kit | 00MW570 | 1 |
| MECHANICAL, SE360V2 SMA Filler | 03LD391 | 1 |
| Mechanical, SE360V2 lock switch kit | 03LC492 | 1 |
| MECHANICAL, SE360V2 CPU air baffle | 03LD392 | 1 |
| Mechanical Assembly, SE360 V2 RISER Cage Assembly (M.2/SSD) | 03KL382 | 1 |
| LABEL,SE360V2 Riser label (7mm + M.2) | 03LF302 | 1 |
| LABEL,SE360V2 Node Label Group | 03LF297 | 1 |
| I/O FILLER Group Bill of Material (GBM) | 03LE692 | 1 |
| Heatsink, M.2 SSD short (non-intel) | 02JJ924 | 1 |
| Front I/O Assembly, Rear, for SE360 V2 with Bluetooth Button (LED), ID Button (LED), Power Button (LED), X86 WiFi Active LED, XCC WIFi Active LED, Activation LED & System Error LED | 03KM989 | 1 |
| Front I/O Assembly, Front, for SE360 V2 with LEDs (Power, ID, BT Enabling, Status, Activation, BMC WiFi Active, x86 WiFi Active) | 03KM986 | 1 |
| Front I/O assembly, Front, For SE350/360 V2 with 4 x 1Gb RJ45 and 2 x 2.5Gb RJ45, 1U IO Board | 03KM993 | 1 |

| | | |
|---|---|---|
| For SE350/ SE360 V2 USB-C to VGA Adapter | 03LF889 | 1 |
| Fan, SE360V2 Fan Assembly (Front to Rear) | 03LA644 | 2 |
| COVER,SE360V2 IO COVER(1G) | 03LD388 | 1 |
| Chassis, SE360V2 Node Chassis with Out Top Cover | 03LA020 | 1 |
| Cable, internal, 370mm, SE360 V2, Front OP Panel to Rear OP Panel, 1X15-1x15 | 03LF068 | 1 |
| Cable, internal, 360mm, SE360 V2, MB to IO board, power, 2x2-2x2 | 03LF080 | 1 |
| Cable, internal, 350mm, SE360 V2, MB to M2 Cabled Adapter, power, 2x2-2x2 | 03LF081 | 1 |
| Cable, internal, 300mm, SE360 V2, IO board to HDD BP1, power, 2x6-2x5 | 03LD306 | 1 |
| Cable, internal, 220mm, SE360 V2, M.2 Cabled Adapter, Gen 4-MCIO 74-74 | 03LF078 | 1 |
| Cable, internal, 205mm, SE360 V2, MB to IO Board, MCIO 148-148 | 03LD634 | 1 |
| Cable, internal, 200mm, SE360 V2, MB to Front OP panel, 2x10-2x10 | 03LF067 | 1 |
| Cable, internal, 180mm, SE360 V2, MB to DC/AC PMB LED, 1x2-1x2 | 03LD288 | 1 |
| Cable, internal, 150mm, SE360 V2, MB to BP2, SATA signal, Swift, SlimSAS 38-38 | 03LF077 | 1 |
| Cable, internal, 130/165mm, SE360 V2, MB to HDD BP1 +BP2, Gen4, MCIO 74-38+38 | 03LF076 | 1 |
| Cable, internal, 100mm, SE360 V2, MB to intrusion switch,1x3-1x2 | 03LC807 | 1 |
| Backplane, SATA/NVMe, 2.5", 1 Bay, Front, Any Bay, 2U 7mm Backplane 2 | 03KM991 | 1 |
| Backplane, SATA/NVMe, 2.5", 1 Bay, Front, Any Bay, 2U 7mm Backplane 1 | 03KM990 | 1 |
| 64GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM | 02JK971 | 4 |
| 2.8m, 10A/120V, C13 to NEMA 5-15P (US) Line Cord | 39M5081 | 3 |

***Lenovo ThinkEdge SE450 Bill of Materials***

| Description | Part Number | Qty |
|---|---|---|
| ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD | 02JG529 | 2 |
| ThinkSystem Intel I350 1GbE RJ45 4-port Internal Lock OCP Ethernet Adapter | 03GX610 | 1 |
| ThinkSystem M.2 7450 PRO 1.92TB Read Intensive NVMe PCIe 4.0 x4 NHS SSD SED (with Heatsink) | 03KH325 | 2 |
| ThinkSystem 1x1 2.5" HDD Filler | 00WF660 | 2 |
| System Label Group | 03GX843 | 1 |
| System Board, Ice Lake, without TPM and Security Pack | 03LF012 | 1 |
| Riser 1, PCIE, x16/x16, slot 2 | 03KL711 | 1 |
| REAR ANTENNA FILLER | 02YF859 | 1 |
| Power Supply Delta 1100W 230Vac/115Vac AC Hot Swap | 02YF647 | 2 |
| Power Middle Plane, SR850 V2 Power Paddle Board | 03LD924 | 1 |
| OCP CABLE BRACKET | 02YF841 | 1 |
| No TPM, Root of Trust Module | 03GX801 | 1 |
| Misc.kit, SCREW | 02YF866 | 1 |
| Misc.kit, LABEL | 02YF843 | 1 |
| Misc. kit, CAP,LABEL | 02YF836 | 1 |
| MECHANICAL, M.2 retainer, Single | 02JJ961 | 2 |
| Mechanical assembly, RISER HDD CAGE | 02YF837 | 1 |

| | | |
|---|---|---|
| Mechanical assembly, REAR TOP COVER 360MM | 02YF869 | 1 |
| Mechanical assembly, PCI RISER BRACKET LEFT | 02YF825 | 1 |
| Mechanical assembly, INTERNAL HDD CAGE | 02YF829 | 1 |
| Mechanical assembly, I/O FILLER | 02YF856 | 1 |
| Mechanical assembly, FRONT TOP COVER | 02YF867 | 1 |
| Mechanical assembly, BEZEL DETECT BRACKET FILLER | 02YF839 | 1 |
| Mechanical assembly, AIR BAFFLE L360 | 02YF832 | 1 |
| Mechanical assembly, 2U SBS M2 BRACKET | 02YF828 | 1 |
| Mechanical assembly, 2U OCP COVER CAGE | 02YF840 | 1 |
| Mechanical assembly, 2U EIA BRACKET RIGHT | 02YF852 | 1 |
| Mechanical assembly, 2U EIA BRACKET LEFT | 02YF851 | 1 |
| Mechanical assembly, 2POST RACK RAIL KIT | 02YF865 | 1 |
| M.2 SATA/NVMe 2-Bay Kit | 01PE603 | 1 |
| LGA4189 1U 3x Heat pipes Standard Heatsink W/Grease | 02YE595 | 1 |
| Label, SYSTEM SERVICE LABEL LANGUAGE INTERNATIONAL | 02YF844 | 1 |
| FAN, 60mm 16K RPM, Cold Swap, Dual Rotor | 03GX838 | 6 |
| CPU,Intel,ICX 5318N 24C/2.2G/165W | 02YH685 | 1 |
| Chassis, 2U,360MM CHASSIS | 02YF823 | 1 |
| Cable, internal, PSU Signal Cable<br>(Pico-Clasp2x15 to Pico-Clasp2x15)<br>(Signal: 140mm | 02YF783 | 1 |
| Cable, internal,OCP Cable<br>(2*MCIO74pin+GenZ 1C to OCP)<br>(MCIO: 220mm/205mm)<br>(GenZ: 50mm) | 02YF774 | 1 |
| Cable, internal, Micro HI 2X8 To Micro-Hi 2x5 245mm | 02YF815 | 1 |
| Cable, internal, MCIO 8x to MCIO 8x<br>184mm | 02YF807 | 1 |
| Cable, internal,M.2 Cable<br>(LP-Slimline4X RA+Pico-Clasp 2x10 to M.2)<br>(190mm/190mm/197mm) | 02YF786 | 1 |
| Cable, internal, Led Cable for Wall Mount<br>125mm | 02JJ722 | 1 |
| Cable, internal,Fan Power Cable with board<br>(Picth2.0 2x4: 70mm) | 02YF787 | 6 |
| Cable, internal, NVMe Cable for 2*7mmNVMe<br>(MCIO8X+Micro-Hi2x8 to NVMe*2)<br>Signal: 260mm/620mm<br>Power: 325mm/360mm | 02YF795 | 1 |
| Cable, internal, LP-Slimline 4x to Slimline 4x<br>260mm | 02YF797 | 1 |

| | | |
|---|---|---|
| BAFFLE FILLER FOR AIR BAFFLE L360 | 02YF834 | 1 |
| Backplane, SATA/SAS/NVMe, 2.5", 2Bay, Front, Any Bay | 03GX832 | 1 |
| 64GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM | 02JK971 | 8 |
| 4U MS PCIe Filler | 01KP742 | 2 |
| 2.8m, 13A/100-250V, C13 to C14 Jumper Cord | 39M5509 | 2 |
| 1U CPU HEATSINK FILLER FOR AIR BAFFLE L360 | 02YF833 | 1 |

# Authors and contributors

We would like to extend our sincere gratitude for their invaluable contribution to the solution design and validation. The insights, expertise, and dedication they brought played a pivotal role in shaping a comprehensive solution. Their collaborative spirit enhanced the quality of both the design and validation processes, ensuring they meet the high standards we strive for. Thank you for the hard work and commitment—it has made a significant impact on the success of this project.

| Name | Team | Title |
|---|---|---|
| Vanita Meyer | Lenovo | Enterprise AI Solution Manager |
| Chris Arenas | Lenovo | Solution Specialist |
| Jay Bryant | Lenovo | LOC-A Consulting Engineer |
| Jean Marius | Lenovo | LOC- A Product Line Manager |
| Pier Ventre | Intel | Cloud Software Dev. Engineer |
| Mark Malinkovich | Intel | Dir Edge Compute & AI Software Product Marketing |
| Ashutosh Kumar | Intel | Software Technical Marketing Engineer |
| Anurag Ranjan | Intel | TIBER Edge Infra Management Product Manager |
| Joy Chua | meldCX | EVP of Strategy & Operations |
| Rachel Melisa | meldCX | Marketing Manager |
| Philnyl Gorial | meldCX | Products Lead |

# Acronyms

| Acronym | Meaning |
|---|---|
| ISV | Independent Software Vendor |
| OS | Operating System |
| GUI | Graphic User Interface |
| AI | Artificial Intelligence |
| LVD | Lenovo Validated Design |
| LOC-A | Lenovo Open Cloud Automation |
| LXCA | Lenovo XClarity Administrator |
| XCC | XClarity Controller (software running as the bare metal management controller) |
| CCTV | Closed-circuit TV |
| CPU | Centra Processing Unit |
| GPU | Graphic Processing Unit |
| RAM | Random Access Memory |
| FPS | Frame per second |
| VM | Virtualized Machine |
| IP Camera | Internet Protocol Camera |
| RTSP | Real-Time Streaming Protocol |
| VPN | Virtual Private Network |
| POS | Retail Point of Sale |
| POE | Power of Ethernet |
| VDI | Virtual Desktop Infrastructure |

# Resources

| Resource | Link |
|---|---|
| ThinkEdge SE350 V2 | ThinkEdge SE350 V2 Datasheet > Lenovo Press |
| ThinkEdge SE360 V2 | Lenovo ThinkEdge SE360 V2 Server Product Guide > Lenovo Press |
| ThinkEdge SE450 | Lenovo ThinkEdge SE450 Edge Server Product Guide > Lenovo Press |
| Intel OpenVINO™ toolkit | Intel OpenVINO™ toolkit |
| Intel Tiber Edge Platform | Intel® Tiber™ Edge Platform |
| meldCX® Website | meldCX® |
| meldCX® Helpdesk | Viana™ by meldCX® helpdesk |
| meldCX® Whitepaper | Whitepaper – Vision Analytics bridging the gap between the physical and digital channels |
| Viana on Intel Partner Showcase | Viana Vision Analytics |
| Lenovo Open Cloud Automation | LOC-A |
| Lenovo XClarity | XClarity System Management |
| Gartner Report | What Edge Computing Means for Infrastructure and Operations Leaders |
| IDC Edge Computing Research | Edge computing lessons learned \| CIO |
| Lenovo + Intel TIBER Edge AI Solution | s7d1.scene7.com/is/content/Lenovoassetsprod/TBT_Len_Int_Strata_091_FILM2_SUBS_V6_VL_TX?refId=bc3f18cd-58dc-42d7-8c6b-7c89252b9060 |

# Document history

| | | |
|---|---|---|
| Version 1.0 | October 10, 2024 | First version includes Lenovo ThinkEdge Systems, Intel® Tiber™ Edge Platform, LOC-A, XClarity and Viana™ by meldCX® |

# Trademarks and special notices



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

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

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

Lenovo®
ThinkEdge®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

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

Azure® and Excel® are trademarks of Microsoft Corporation in the United States, other countries, or both.

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

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

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