Redefining distributed applications and edge AI

with seamless solutions built on Intel® Tiber™ Edge Platform and edge-optimized Lenovo infrastructure

Edge AI offers powerful new opportunities to innovate, but the complex nature of deploying, managing, and integrating AI systems at the edge is holding many organizations back.

To help solve these challenges, Lenovo and Intel® have created a new generation of edge AI solutions built on Intel® Tiber™ Edge Platform, Lenovo Open Cloud Automation, and the Lenovo XClarity family of management tools. These technologies are seamlessly integrated in the edge-optimized Lenovo ThinkEdge Series of devices.

Intel® Tiber™ Edge Platform is a new, commercial software platform that enables enterprises to build, deploy, run, manage, and scale edge and AI solutions on standard hardware with cloud-like simplicity. The platform provides simple and intuitive control of the automations and management solutions built into ThinkEdge devices, including Lenovo Open Cloud Automation (LOC-A) and Lenovo XClarity. LOC-A enables organizations to rapidly deploy edge AI infrastructure at scale, in multiple locations and with just a few touches. Lenovo XClarity provides deep remote visibility into Lenovo hardware at the edge. And with the OpenVINO™ toolkit, developers can access built-in AI inferencing capabilities that are optimized for the low latency, low power characteristics of edge AI infrastructure.

Learn more about how the seamless integration between Intel® Tiber™ Edge Platform and the Lenovo ThinkEdge portfolio is redefining edge AI.

Addressing obstacles to edge AI innovation

More enterprises are leveraging edge AI than ever before, in their operations on factory floors, in restaurants, in logistics facilities, and at other locations. According to a 2023 research report by Accenture:

- 83% of believe edge computing will be essential to remaining competitive in the future.
- 98% of global executives say AI foundation models will play an important role in their organizations' strategies in the next three to five years.
- The most advanced edge adopters are more innovative, more efficient, and run more cost-effective operations than their competitors.

Edge AI capabilities are therefore becoming key to competitiveness. With edge AI, insight is IT is no longer confined to standalone data centers. It is instead distributed across a mix of on-premises servers, clouds, and growing arrays of infrastructure at the edge. Organizations can use this agile IT to create faster and smarter business insights.

These insights can promote live interventions such as defect identification, hazard detection, and predictive machine maintenance, elevating processes to unprecedented levels of efficiency. And capabilities such as real-time optimization and cognitive computing help to improve operational safety, boost productivity, and provide workforces with actionable insights for decision making. In a manufacturing setting, for example, edge AI brings cost optimization, enhanced product quality, and advanced inventory management to the fore. And by reducing cloud traffic and energy consumption, edge AI is carving a pathway to sustainability — a significant step toward a greener future.

Yet only 65% of companies are using edge today, according to Accenture.[1] Many are still grappling with the many challenges of implementing and managing edge AI solutions, which include:

- **Complexity of integrating heterogenous systems** – Edge AI solutions must integrate new and existing heterogeneous infrastructure and applications, because replacing existing mission-critical edge systems is often not an option.
- **Difficulty of running AI workloads efficiently at the edge** – Effective edge AI solutions must typically maintain low latency performance and high energy efficiency. Achieving this requires skilled AI development and optimization, as well as the use cost-effective, edge-optimized hardware.
- **Managing remote and distributed infrastructure** - Deploying, running, and managing edge AI solutions is extremely time-consuming as they may encompass hundreds or thousands of distributed nodes and applications. Solutions may also extend across diverse sites and geographies, with varied regulatory requirements.

Edge and AI solution development can be achieved through multiple paths and may require a mix of various technologies, products, and partners. Analysts Gartner suggest that "infrastructure and operations leaders need to evaluate solutions that accelerate deployments and support extensibility." But if organizations had a foundation for edge AI that could simplify the tasks of integrating diverse systems, optimizing AI workloads, and managing distributed infrastructure, the path could be made easier and less risky.

Many edge and AI use cases across industry verticals

Industrial facilities

Identifying product defects
Predictive maintenance with computer vision
Process optimization
Predicting demand and adjusting production

Telecommunications

Private 5G networks
Remote network infrastructure
Network service optimization
Efficient energy management

Smart cities and transportation

Traffic management
Predictive maintenance
Improving public safety
Autonomous vehicles

Retail and professional services

Understand consumer demands
In-store traffic analytics
Automated checkout
Smart shelving

Bringing cloud-like simplicity to edge AI

Recognizing these challenges, Lenovo and Intel® have partnered to offer comprehensive solutions designed to simplify the architecture, deployment, optimization, and remote management of edge AI infrastructure. They offer seamlessly integrated technologies, thanks to Lenovo and Intel's long-standing partnership and co-engineering initiatives.

At the heart of these solutions is the **Lenovo ThinkEdge Series** of edge-optimized servers. Compact, resilient, and energy-efficient, ThinkEdge devices deliver the power of Intel® processors to run AI workloads of all sizes in the most challenging of locations. ThinkEdge devices are also embedded with the **Lenovo XClarity Controller (XCC)**, the next-generation server and baseboard management controller (BMC) that enables deep observability into server health and configuration, as well as remote provisioning and management. Together with **Lenovo XClarity Administrator**, XCC can automate discovery, inventory, tracking, updates, monitoring, and provisioning for a wide range of Lenovo systems including servers, storage, and network switches.

At the software layer, XClarity Controller integrates with both **Lenovo Open Cloud Automation (LOC-A)** and **Intel® Tiber™ Edge Platform** to completely transform the way edge AI infrastructure is deployed, managed, and optimized. Intel® Tiber™ Edge Platform is a new, edge-native commercial software platform that provides a single plane of control, enabling enterprises to build, deploy, run, manage, and scale their edge AI solutions with cloud-like simplicity. And Lenovo Open Cloud Automation (LOC-A) provides powerful automations for the deployment and lifecycle management of hundreds or thousands of edge nodes, making it easy to onboard devices and run edge AI infrastructure at scale.

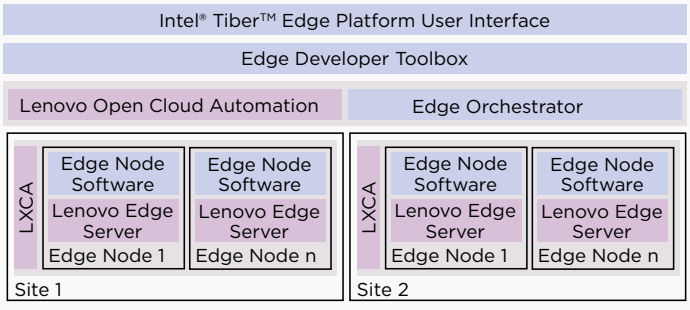

Figure 1: Intel® Tiber™ Edge Platform provides a single-pane of control for distributed edge AI infrastructure including Lenovo LOC-A and Lenovo XClarity deployment and management capabilities

Together, Intel® Tiber™ Edge Platform and the Lenovo ThinkEdge portfolio deliver a seamless experience combining truly edge-native capabilities for security and near-zero-touch management, built on our deep industry experience and unrivaled ecosystem.

Meet diverse needs with a broad ecosystem

Lenovo and Intel® solutions are built on our combined expertise and help solve the challenge of how to manage diverse, distributed infrastructure.

- **Proven edge expertise** – Intel® Tiber™ Edge Platform leverages the company's experience from supporting more than 90,000 edge deployments today, with a footprint spanning more than 200 million processors sold in the past ten years. The Lenovo ThinkEdge Series offers a long-established portfolio of purpose-built edge infrastructure.
- **Support for diverse hardware** – Seamlessly integrate almost all Lenovo servers, storage, networking switches, and more in one easier-to-manage solution.
- **Leverage heterogeneous components** – Integrate heterogeneous and existing brownfield components into edge solutions.

Optimize AI performance at the edge

Developers can use built-in tools for AI model development and management, plus AI inference runtime optimizations. These include:

- The AI Developer toolbox with components for building AI and applications for edge and hybrid implementations.
- Intel® Geti™, a software platform for building computer vision models with less data and time.
- The OpenVINO™ toolkit, specially optimized for low latency, low power, and existing hardware at the edge, enabling standard hardware already deployed to run AI applications efficiently.
- Granular application orchestration capabilities to target existing workloads, AI initiatives, industrial controls, networking, and more, helping them to run most efficiently for seamless use case delivery.

Deploy industry-ready solutions

With Lenovo and Intel® solutions, organizations can access a huge range of edge AI solutions that have been proven to solve real-world business problems.

- Lenovo offers more than 150 turnkey AI solutions for diverse use cases, built on proven software from the Lenovo AI Innovators ecosystem of independent software vendor (ISV) partners.
- The Intel® Edge Insights System helps organizations manage, operate, and deploy AI at the Edge. Integrated data management and ingestion for time series and video data enables industrial use cases such as defect detection and predictive analytics for manufacturing, energy, and beyond.
- Intel® SceneScape can track and manage physical environments using digital twin capabilities for enhanced spatial awareness.

Deep observability with ThinkEdge and Intel® Tiber™ Edge Platform

The integration between LOC-A, Intel® Tiber™ Edge Platform, and Lenovo XClarity makes it easy to onboard new edge nodes, create and manage clusters, and gain deep observability into each node, all in a single pane. Organizations can:

- **Onboard bare-metal edge nodes faster** and with fewer interactions.
- **Deploy the operating system remotely,** even in network-constrained environments.
- **Create clusters that include newly onboarded edge nodes,** with a near-zero-touch cluster orchestration capability, and then view or edit cluster properties.
- **Easily browse, deploy, modify, and upgrade new applications** to a cluster using the application catalogue.

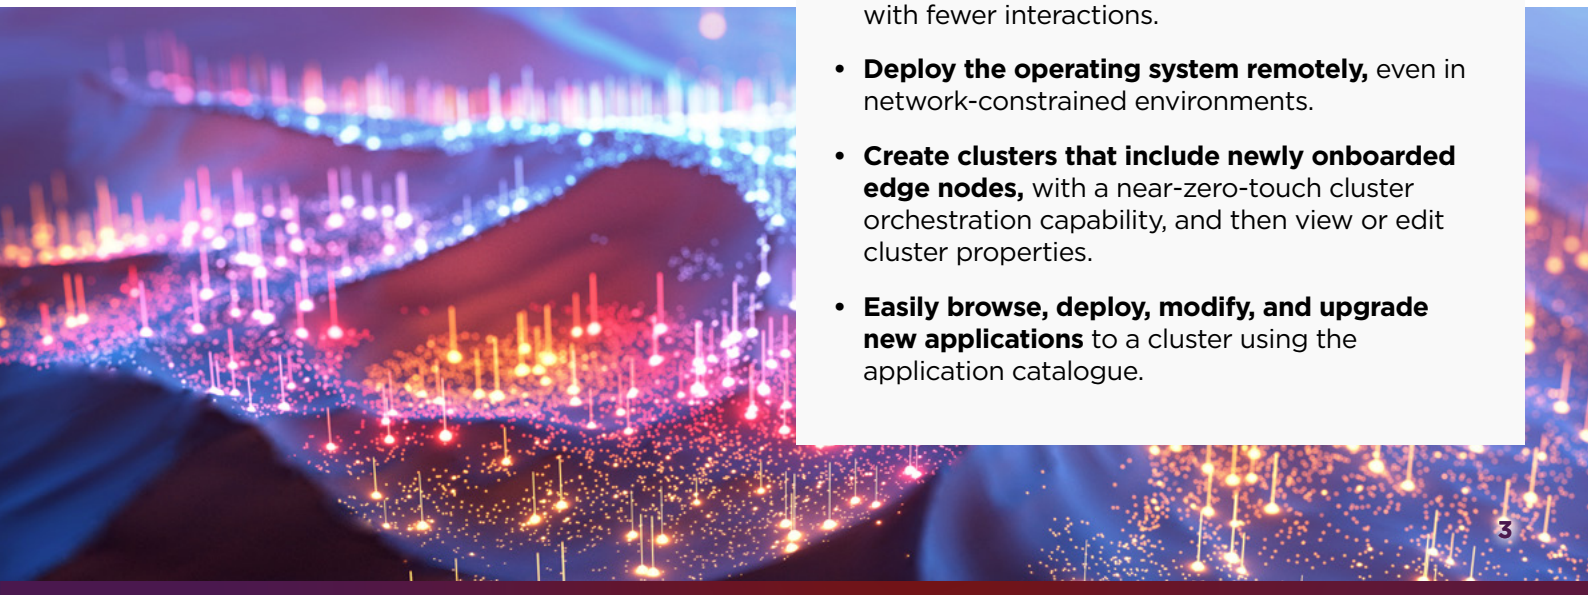

Seamlessly integrating edge AI innovations

Lenovo and Intel® edge AI solutions bring together a wide range of technologies to deliver seamless and simplified infrastructure deployment, management, and performance optimizations.

Build in edge expertise with Intel® Tiber™ Edge Platform

Intel® Tiber™ Edge Platform is an edge-native commercial software platform that enables organizations to build, deploy, run, manage, and scale edge and AI solutions on standard hardware with cloud-like simplicity. Lenovo hardware is engineered to take full advantage of its capabilities.

The benefits of Intel® Tiber™ Edge Platform include:

- **Secure onboarding and management** of a fleet of edge nodes.
- **Support for diverse components** including a range of architectures, accelerators, and third-party application.
- **Performance tuning for AI workloads and applications** on a wide spectrum of edge hardware, with built-in AI inference enabled by OpenVINO™ runtime optimizations.
- **Easier Day 0/1/2 operations** across an edge fleet.

Key features of Intel® Tiber™ Edge Platform

Single-pane-of-glass management

Enable seamless infrastructure and application administration across geographies and workflows.

AI and application development

Use powerful low-code to high-code AI model/app development and optimization, plus digital twin capabilities to deploy full-stack enterprise AI applications.

Built-in OpenVINO™ AI inference run-time optimizations

Achieve desired speed, accuracy, and power efficiency on right-sized components.

Deep hardware-aware telemetry

Track and automate application deployment based on policies and observability.

Zero-touch lifecycle management and orchestration

Unify management of applications, infrastructure, and AI across edge locations for Day 0/1/2 operations through policies, enabling teams to focus on business logic. full-stack enterprise AI applications.

Dynamic policy-based workload placement

Tackle connectivity challenges and workload efficiencies between near edge (remote data servers) and far edge (HMIs, gateways) infrastructure.

Zero trust security

Features include safe device onboarding/boot, protected communications, full disk encryption, identity, and access management.

Automate edge deployments with LOC-A

Lenovo Open Cloud Automation (LOC-A) is a software solution that rapidly deploys and manages the lifecycle of server infrastructure at scale, however remote and numerous your edge locations may be. The solution uses Infrastructure as Code and GitOps best practices for consistent and error-free results. This speeds up on-premises cloud deployment on bare-metal servers and helps maintain a good security posture by avoiding configuration drift.

When it comes to deployment options, LOC-A adapts to each environment and can be installed on bare-metal, containers or virtual machines. Regardless of how it is installed, LOC-A manages the lifecycle of on-prem clouds from Red Hat, VMware, or Kubernetes. At the edge, LOC-A solves the logistical challenge: with zero-touch-provisioning of Lenovo edge servers and remote deployment of edge cloud software, bare-metal servers can be shipped directly to site, eliminating the need for golden images and staging environments.

Maximize time to value with low touch deployment.

- Create repeatable deployment templates and configurations that can be used to accelerate distributed edge deployments.
- Deploy distributed Edge infrastructure up to 3x[2] faster through automation instead of manual processes.
- Reduce number of resources by 4x[2] and CO2 emissions by 65%.[2]

Simplify operations and reduce costs.

- Have edge infrastructure ready for use in minutes, not days or weeks. Operational automation framework reduces human errors, while making it easy to keep up with latest software and updates.
- Easily scale across multiple locations and clusters.

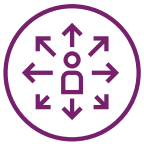

Efficiently extend networks with an Edge friendly solution.

- Reduce truck rolls for edge deployments, as automation becomes the new baseline for expansion.

Manage hardware like software with Lenovo XClarity tools

The Lenovo XClarity family of systems management tools and software helps organizations to further improve efficiency, reduce costs and enhance the availability of edge AI infrastructure. Lenovo XClarity manages Lenovo ThinkEdge infrastructure, integrates with the LOC-A software platform, and Intel® Tiber™ Edge Platform, making it even easier to deploy and manage distributed edge infrastructure from a single intuitive interface.

Lenovo XClarity family tools

Lenovo XClarity Administrator (LXCA)

A centralized resource management solution aimed at reducing complexity, speeding response, and enhancing availability of infrastructure. LXCA runs as a virtual appliance and provides agent-free hardware management that automates discovery, inventory, tracking, updates, monitoring, and provisioning for Lenovo ThinkEdge systems, storage, network switches, hyperconverged and ThinkAgile solutions.

Lenovo XClarity Controller (XCC)

An integrated service processor that provides advanced service-processor control, monitoring, and alerting functions, accessible via CLI or web interface and integrated with LXCA and LOC-A.

Manage hardware like software with Lenovo XClarity tools

The Lenovo XClarity family of systems management tools and software helps organizations to further improve efficiency, reduce costs and enhance the availability of edge AI infrastructure. Lenovo XClarity manages Lenovo ThinkEdge infrastructure, integrates with the LOC-A software platform, and Intel® Tiber™ Edge Platform, making it even easier to deploy and manage distributed edge infrastructure from a single intuitive interface.

Lenovo ThinkEdge Series devices

ThinkEdge SE10

- Powerful and cost-effective entry-level edge solution.
- Easily configured to adapt to a broad range of edge vertical applications with 24/7 serviceability.

Lenovo ThinkEdge SE360 V2

- Designed and built for the unique requirements of edge AI applications and inference.
- Rugged, compact-sized solution for harsh environments.

Lenovo ThinkEdge SE450

- Powerful edge AI server featuring Intel® Xeon® processors and up to 1 TB memory.
- Broad connectivity and reliable operation in extreme temperatures.

Make AI more accessible with OpenVINO™ toolkit and Intel® Geti™

Solutions are also ready to run the Intel distribution of the OpenVINO™ toolkit, which helps to make generative AI more accessible for real-world scenarios. The toolkit helps to streamline AI development and integration of deep learning in domains like generative AI, computer vision, and large language models. Developers can accelerate AI inference with lower latency and higher throughput while maintaining accuracy, reducing model footprint, and optimizing hardware use.

The OpenVINO™ toolkit provides a three-step workflow for AI development:

1. **Model** – Train models with popular frameworks like TensorFlow and PyTorch. Access 280+ open source and pre-trained models with Open Model Zoo. And use Intel® Geti™, a new software platform, to build computer vision models in a fraction of the time and with less data.

2. **Optimize** – Convert and optimize models with tools including Model Convertor for OpenVINO™ and Direct Model Conversion for TensorFlow and PyTorch.

3. **Deploy** – Across your Lenovo edge AI infrastructure environments. Serve models over cGRP, REST, and C API endpoints. And inference at the edge with OpenVINO™ Runtime, which offers multiple inference modes to allow optimum hardware utilization under your specific conditions.

Learn more

Edge AI offers innovation and competitive advantage to organizations in a wide range of industries, but the complexity of deploying and managing infrastructure has proved too large an obstacle for some organizations.

Lenovo and Intel solutions provide the ideal path for organizations who want to take full advantage of Intel® Tiber™ Edge Platform, powerful automations for edge AI deployment and lifecycle management with LOC-A and Lenovo XClarity, and efficient ThinkEdge servers. Organizations can integrate heterogeneous and existing brownfield components into edge solutions, start developing AI models and applications right away, and easily manage the full lifecycle of distributed infrastructure from a single pane of glass.

For more information, contact your Lenovo representative.

Lenovo Open Cloud Automation **Intel® Tiber™ Edge Platform**

intel. + Lenovo

[1] Accenture (2023). Leading with edge: How to reinvent with data and AI. https://www.accenture.com/us-en/insights/cloud/edge-computing.
[2] Gartner (2023, October 18th). Predicts 2024: Edge Computing Technologies Are Gaining Traction and Maturity. https://www.gartner.com/en/documents/4850031.