invisible ai

cisco

Invisible AI and Cisco Accelerating AI-Driven Machine Visioning

May 2026

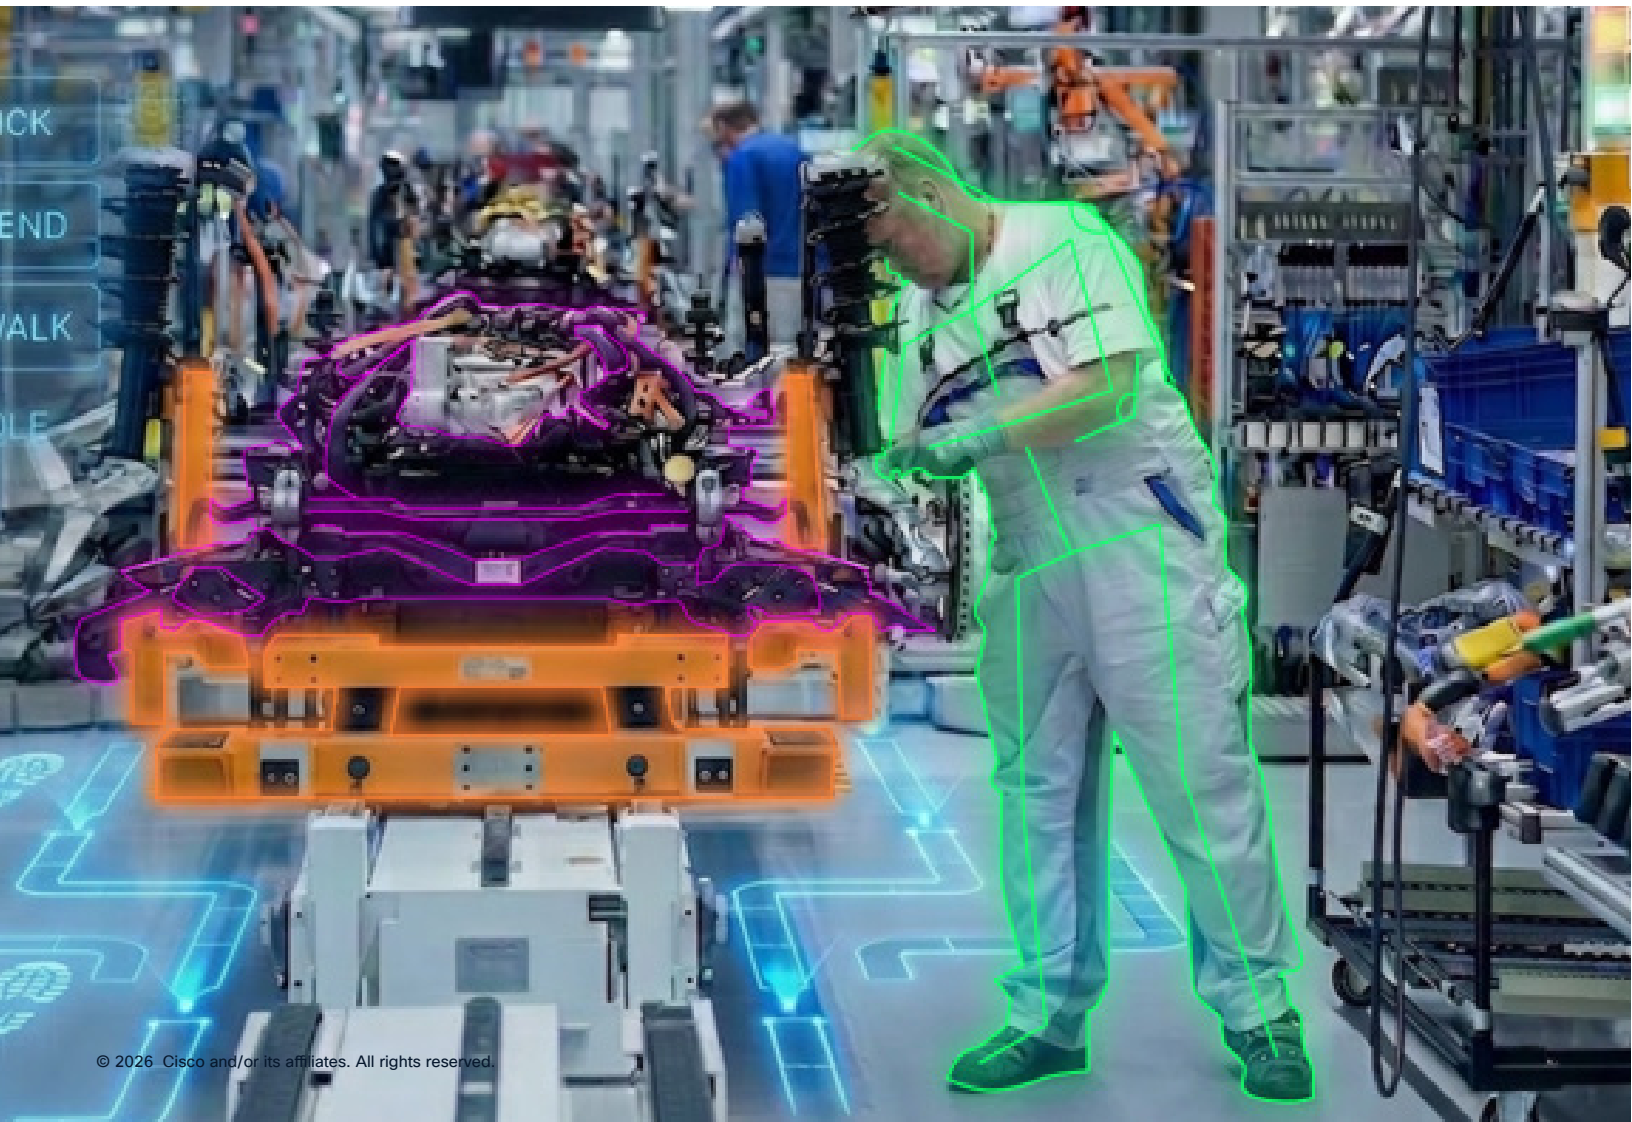

invisible ai

cisco

Contents

Overview and business drivers for AI-driven machine vision applications

Overview

Modern operational excellence is about more than just hardware and robotics; it is increasingly driven by advanced computer vision systems that transform raw visual data into operational intelligence.

Artificial Intelligence (AI) and Machine Learning (ML) are no longer confined to data centers or cloud applications. They are moving right onto the production line. By integrating AI-driven machine vision, manufacturers are achieving levels of precision, speed, and safety that were previously unthinkable. However, there is a critical, often overlooked component to this transformation: the network.

If AI is the brain of the modern factory, the network is the nervous system. Without a high-performance, resilient, and secure network, even the most sophisticated AI models will fail to deliver. Additionally, a well-designed network can accelerate business outcomes and lower costs by simplifying the deployment of AI-driven vision systems.

This white paper summarizes the collaboration between Cisco and Invisible AI. It presents a description of the business drivers and challenges, plus an overview of the solution. Our collaboration helps customers, machine builders, system implementers, and suppliers accelerate their adoption of next-generation machine vision capabilities.

Target audience

This white paper is for CIOs, CTOs, IT/OT, operational excellence and plant leaders who are scaling AI-driven machine vision in industrial environments and need a repeatable, secure, operationally resilient blueprint. It provides AI shop-floor tools for frontline workers. It is also relevant to manufacturing executives sponsoring digital transformation initiatives focused on throughput, quality, safety, and labor productivity—where faster deployment, lower operational risk, and enterprise-grade security controls are prerequisites for scaling from pilot to multisite rollout.

Business drivers

For operations decision makers, the adoption of AI-driven machine vision isn't just about upgrading technology; it's about driving tangible business outcomes. The benefits are wide-reaching:

Improve Overall Equipment Effectiveness (OEE): Close the operations gap between actual and expected output with actual data, while increasing productivity, reducing waste, improving uptime through real-time AI-analysis of constant visual inspection of cell and line level processes.

Achieve evidence-based problem resolution: Improve root cause analysis of any missed operational targets or downtime with structured analysis of what really happened.

Increase safety: Analyze and optimize worker integration into automated operations, not just to identify the root cause of any incidents, but to proactively identify safety improvements.

Critical challenges facing customers deploying AI-driven machine vision systems

While the benefits of AI-driven machine vision systems are clear, the path to implementation is fraught with challenges. Machine vision systems are not typical in IT applications; they operate in the harsh, high-stakes environment of cell/area zones in production environments.

- Average **cost** of Ethernet cable drops in North America is $2500/per location. Customers want to consolidate data, power, and synchronization onto one cable to optimize the drop locations they already have and use the connections more effectively.
- Machine vision **data access** is critical as images often need to be preserved for compliance and historical reasons and is needed to improve inferencing algorithms—the network needs to support not just the processing at the edge but the secure transfer of vision data to the plant data center and cloud.
- Customers worry that the **amount and size of the vision data** communications will overwhelm and impact existing industrial automation traffic, driving up latency, jitter, and packet loss and leading to outages and downtime.
- **Cybersecurity** is a growing concern, and the vision system data often represents significant intellectual property. Enterprise-level customers want to know how to integrate network and IT security while also protecting the vision systems and the sensitive data they produce.

Our collaboration

Introducing Invisible AI

Invisible AI is an AI-driven machine vision platform that converts video from existing industrial environments into actionable operational intelligence—without requiring changes to standard work or disrupting production. The platform analyzes how work, automated and human, is actually performed and identifies the drivers of lost productivity, quality escapes, and safety risk using computer vision models tuned for industrial operations.

Invisible AI's solution is an edge-first architecture. Video is processed at or near the line to minimize bandwidth burden and keep sensitive visual data close to the source. Invisible AI's Edge Data Pipelines, which perform real-time inference on NVIDIA silicon, require an industrial network to match their demands for power, connectivity, and synchronization. Cisco's high-wattage Power over Ethernet (PoE) (up to 90W) and high-speed industrial switches are recommended edge network infrastructure that helps ensure that the edge processing node has the power and bandwidth to operate without latency, fulfilling the "edge-first architecture" promise.

Key Invisible AI features include:

- **Actionable insights, not raw video:** Invisible AI's technology turns video into structured operational signals—cycle time, delays, process variation, and event markers—so teams can prioritize the highest-impact improvement opportunities.
- **Fast root-cause identification and continuous improvement:** Engineers and operations teams can quickly find "what happened, when, and why," accelerating time to resolution for recurring issues and enabling sustained performance gains.

- **Enterprise integration readiness:** Insights can be operationalized via dashboards and integrations into existing reporting, continuous improvement, and incident workflows (based on each customer's standards and systems).
- **Insights provided in plain terms:** Invisible AI puts problems and solutions in plain terms for executives and supervisors driving decisions that help plants run better by making process performance, variations, and abnormalities visible and measurable at scale.

Better together: Invisible AI and Cisco—Why this is peanut butter + chocolate

Invisible AI makes manufacturing performance visible and speeds up the operation from the floor to the office. Cisco makes it deployable, manageable, and scalable. Together, you get AI value without enterprise risk—and a repeatable path from pilot to multisite rollout based on our collaboration and Cisco's Machine Vision Validated Design.

AI-driven machine vision usually stalls in the middle:

- It starts as an operations win, then hits **network complexity, security reviews**, and **scale pain**.
- Plants worry that vision traffic will **overwhelm industrial automation networks** (latency, jitter, and packet loss = outages).
- Security teams worry that the vision system and video data represent **high-value intellectual property** and a larger OT attack surface.
- Financial costs of physical deployment are real (e.g., expensive drops and infrastructure constraints).

Invisible AI is the "why"—it turns video into operational intelligence teams can act on. **Cisco** is the "how"—it provides the industrial networking and security foundation to run that workload reliably and safely at scale (industrial switches + visibility + segmentation + zero-trust remote access + centralized management). So instead of "a cool pilot that can't scale," **our collaboration enables enterprise capability with a playbook**—standardized architecture, predictable performance, and security by design.

Invisible AI brings these outcomes:

- Creates a "digital evidence layer," converting factory operations and work in real time to structure data for each and every cycle. This structured data feeds the intelligence pipeline that leaders can trust to drive action, not anecdotes.
- Turns video into **measurable operational signals** (performance loss, variation, abnormal conditions) for faster root cause analysis and sustained improvement.
- Establishes a novel two-layer AI system that creates data as the source and applies agentic AI tools to consume and review the situation in accordance with the engineering workflows, essentially enabling every cycle on every station to be observed by an industrial engineer, quality engineer, or supervisor while empowering front-line teams with the latest AI tooling to improve the operational performance at the source.

Cisco brings the confidence to scale:

- **Industrial-grade connectivity plus PoE**, so cameras and edge nodes deploy faster and more cleanly.
- **Continuous OT visibility** (what's connected, what it's talking to, what changed) to reduce blind spots.
- **Segmentation and policy** so vision workloads don't sprawl across the OT environment.
- **Zero-trust remote access** purpose-built for OT (no "forever VPNs") for support and operations.
- **Centralized management** to run this like infrastructure, not like a science project.

Together our collaboration, especially on the Cisco® Validated Design drives:

- **Faster deployment** with a standard design, leading to fewer one-off exceptions.
- **Lower operational risk** by relying on tested performance with design recommendations.
- **Shorter security approval cycles** (clear controls, visibility, governance).
- **Repeatable multiplant scaling** with a validated architecture.

Case study

Hitting target TAKT in hours—Automotive supplier improves throughput by 63% on day 1

Customer profile

An enterprise automotive supplier facing rapid expansion of demand and the need to increase line output quickly, in other words improve the TAKT (or cycle) time.

The challenge

During a surge in vehicle demand, the supplier needed to raise throughput from 320 to 520 parts per shift, but the line repeatedly fell short. While the team suspected bottlenecks, they couldn't pinpoint where interruptions were happening or why they were occurring, making countermeasures slow and difficult to validate.

The solution

Invisible AI deployed AI edge devices to create a digital twin of the production process and captured real-time metrics tied to bottlenecks and performance losses. Within hours, the team used video-backed insights to focus improvement actions on three high-impact areas:

- **Operator performance coaching:** Video playback of the ideal process helped accelerate execution where speed losses were concentrated.
- **Line rebalancing:** Work was shifted off an overloaded station to reduce downstream starvation and interruptions.
- **Part layout improvements:** Walking and motion waste were reduced by adjusting part presentation and placement.

This approach aligns with a modern edge-first manufacturing architecture: fast deployment, on-premises processing, and operational feedback loops that don't depend on cloud latency.

Results (day 1 impact)

- Throughput increased by 63% per shift on day 1.
- The supplier reported achieving its target TAKT time within 4 hours of deployment.
- As the customer summarized, "The fact that Invisible AI was able to help us meet our target TAKT time within 4 hours of deployment is truly remarkable."

Why it matters

This case demonstrates how secure, scalable edge deployments can convert "known bottlenecks" into measurable, verifiable fixes—quickly. When operational video intelligence and real-time line metrics are captured reliably at the edge and made accessible to Continuous Improvement (CI) leaders, manufacturers can move from "continuous improvement projects" to continuous improvement execution, with outcomes visible in the first shift.

Key network requirements for AI-driven machine vision bandwidth

Machine vision can introduce sustained, high-throughput traffic driven by camera count, resolution, frame rate, and encoding settings. Successful deployments plan bandwidth headroom and avoid sharing constrained links with latency-sensitive automation traffic. A scalable architecture supports growth from pilot (a few cameras) to production rollouts (dozens to hundreds of cameras per facility) without rework.

- **End-to-end connectivity:** Invisible AI's dual-AI model requires structured data transfers between the AI-enabled cameras and AI-enabled servers to process line- or cell-level analytics. Additionally, anomalies or identified issues are typically collected and stored for compliance and historical reasons. And often the raw or processed data from the Invisible AI cameras may be used to improve inferencing engines produced by cloud-based AI engines.
- **PoE:** PoE simplifies camera deployment by reducing electrical work, shortening installation timelines, and improving placement flexibility. High-wattage PoE (90W 4PPoE) options enable broader camera and sensor coverage—including more demanding devices—without redesigning power distribution.
- **Synchronization:** Accurate time synchronization is essential for correlating events across cameras, lines, and systems (e.g., production events, downtime logs, quality incidents). A reference design should support consistent time services so operational signals and evidence align reliably across the environment.
- **Prioritization:** Machine vision traffic must be managed so it does not overwhelm OT networks or create performance risk for automation systems. Quality of Service (QoS) and traffic policy help prioritize critical industrial automation communications while still supporting robust vision workloads—especially during peak load or when troubleshooting.
- **Security:** Video and derived operational data can represent sensitive intellectual property and may be subject to strict enterprise controls. Successful rollouts implement segmentation, least-privilege access, continuous asset visibility, and secure remote operations—without relying on ad hoc exceptions that slow approvals or expand the attack surface.

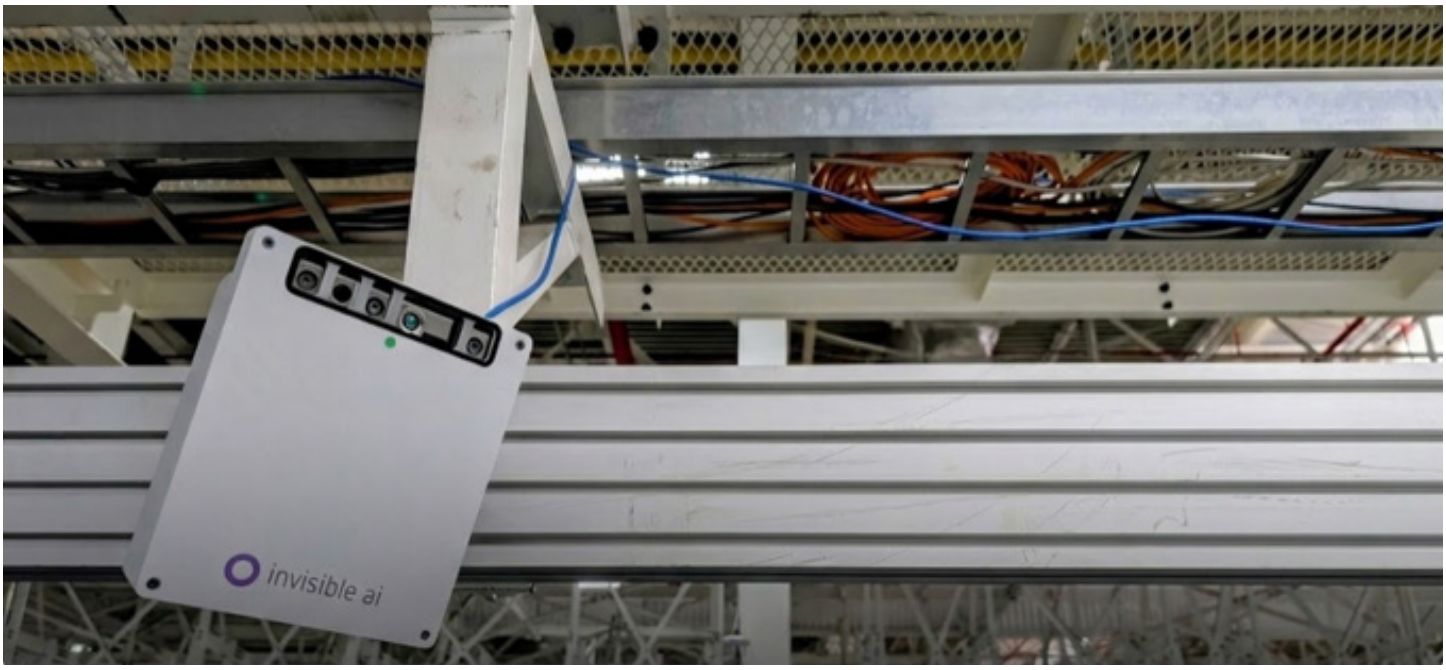

Machine vision reference architecture

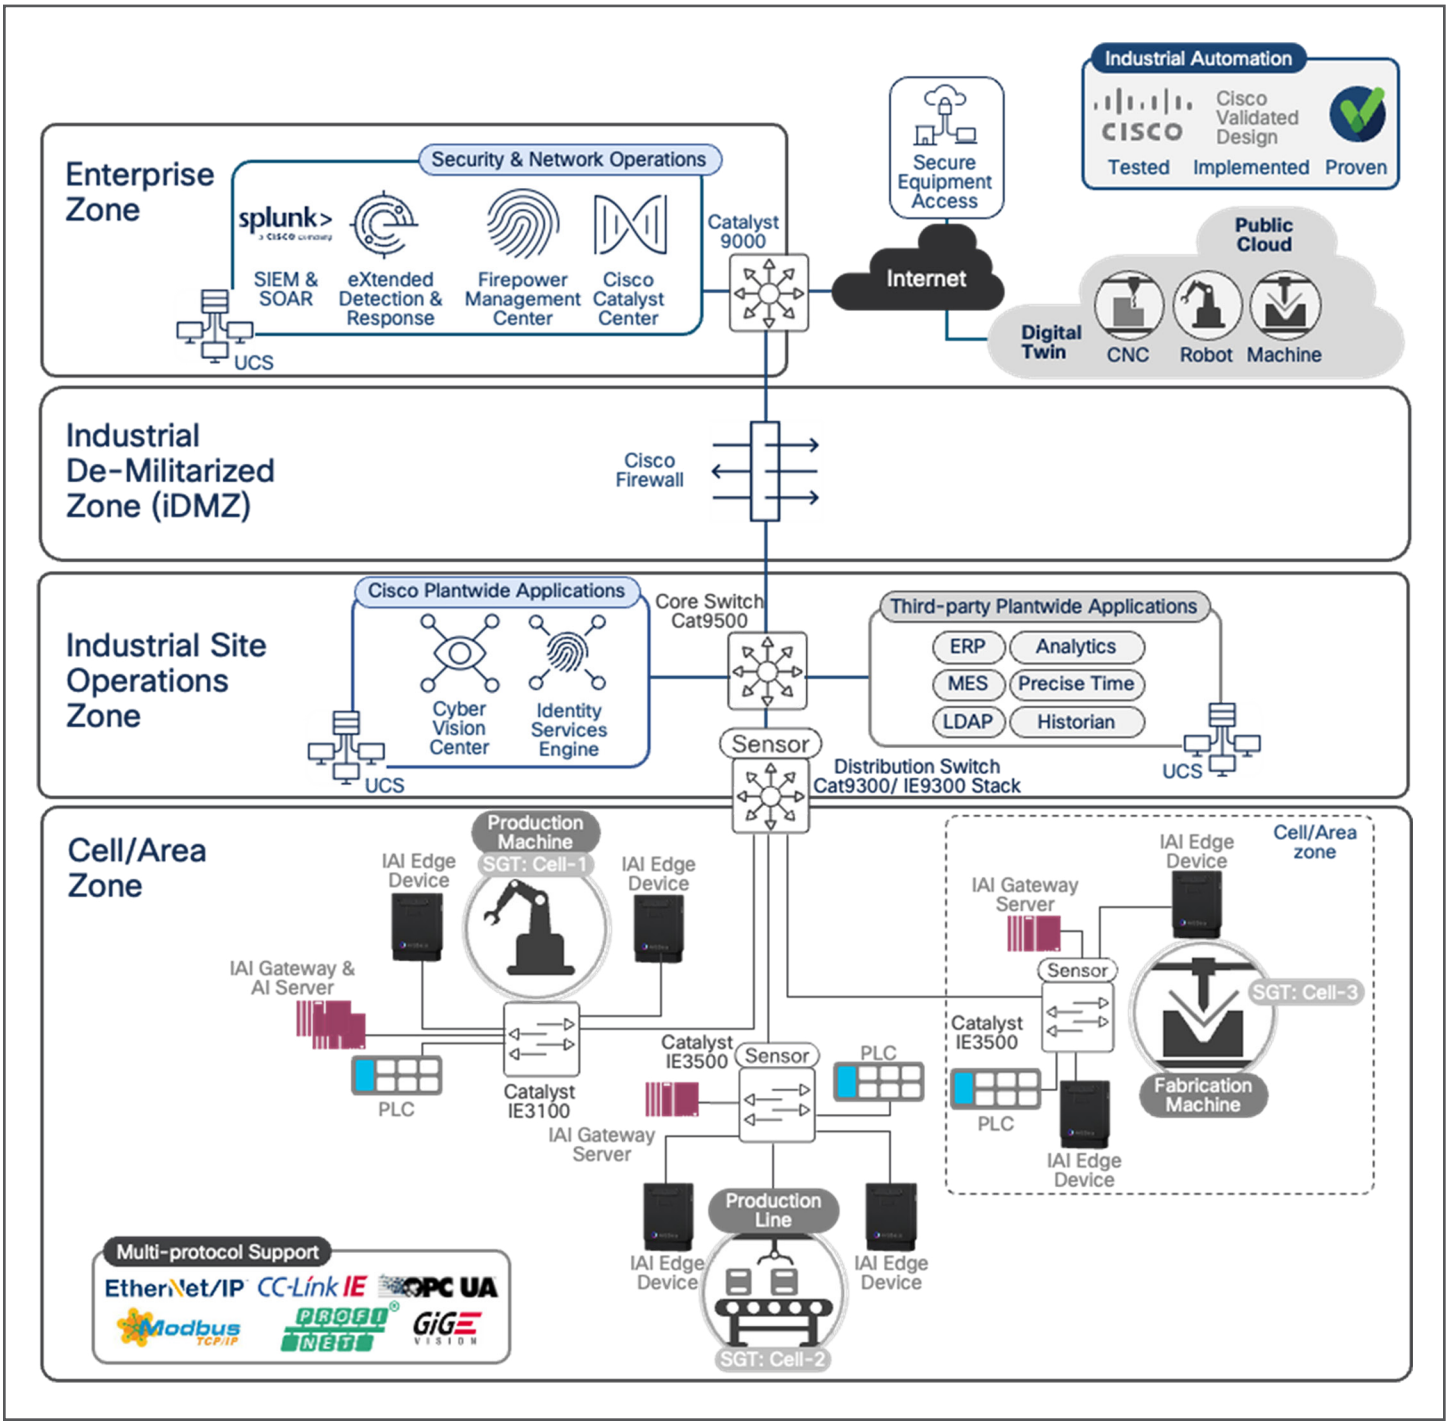

Figure 1. Invisible AI and Cisco machine vision network architecture

Key components

Invisible AI components

Invisible AI machine vision solution components typically include:

- **Industrial 3D edge devices (PoE capable):** Deployed across targeted processes to capture operational activity with flexible placement and minimal facility disruption. An edge inference/processing node on these devices performs video ingestion and AI analysis close to the line to reduce network strain and enable scalable deployments.
- **Orchestration intelligence layer (gateway server):** Transforms video into structured events and metrics (e.g., cycle time variation, delays, micro-stoppages, process adherence signals) so teams can prioritize improvement actions; also serves up the GUI/application for end users.
- **AI server (NVIDIA Spark):** Invisible AI offers custom on-premises LLM and VLM capabilities in partnership with NVIDIA to deliver unprecedented insights for the shop floor.
- **Investigation and collaboration workflows:** Enables teams to quickly locate relevant time windows, correlate operational context, and accelerate root-cause analysis and continuous improvement.
- **Enterprise management and integration interfaces:** Supports fleet-style management and optional integrations into existing IT/OT systems and reporting workflows.

Figure 2. Invisible AI machine vision solution components

Cisco components

Cisco Industrial Ethernet (IE) switches: Available in DIN-rail, IP67-rated, and rack-mount form factors, the Cisco Catalyst™ IE9300, IE3500, and IE3100 Rugged Series and IE3500 Heavy Duty Series switches deliver high-speed ports including 1GE, 2.5GE, and 10GE options, to support demanding industrial AI applications. The switches include high-wattage PoE (up to 90W per port, up to 720W total per switch), enabling flexible deployment of cameras and sensors. They integrate robust security with Cisco Trust Anchor, Cisco Cyber Vision, Cisco TrustSec® segmentation, and Cisco Secure Equipment Access for zero-trust remote management of connected assets. These switches offer high-capacity, low-latency Layer 2 and 3 switching, all managed by Cisco Catalyst Center.

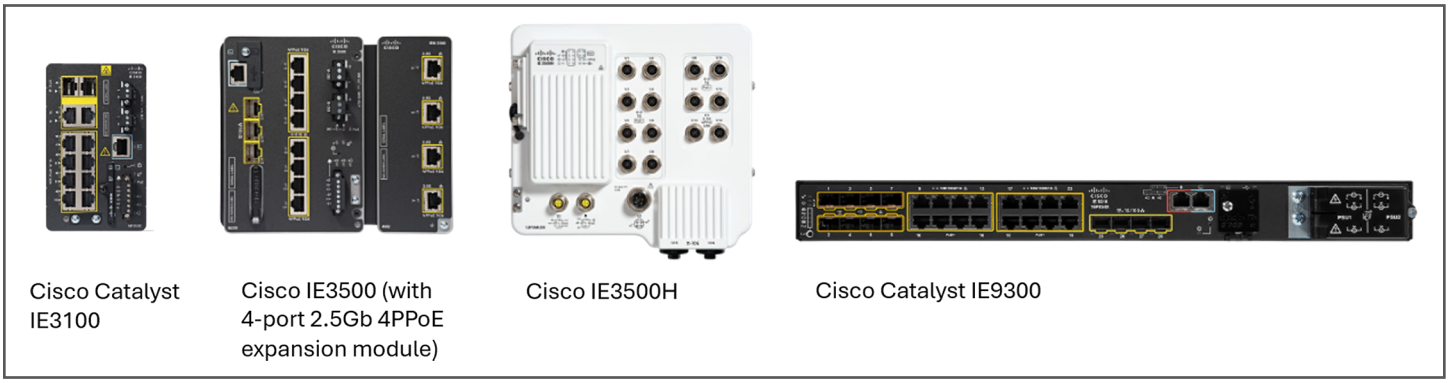

Figure 3. Cisco Industrial Ethernet switches

Cisco Cyber Vision provides deep, continuous visibility into connected OT assets by automatically inventorying devices and mapping their communication patterns. It assesses vulnerabilities by highlighting those being actively exploited, enabling prioritized remediation. The solution continually monitors for threats, helping with early detection and response. Cyber Vision leverages AI to group assets based on their industrial roles, helping operations teams define logical zones for network segmentation policies. Integrated within Cisco industrial switches, this combination delivers a unified networking and security architecture that simplifies deployment, enhances operational resilience, and enables adaptive segmentation aligned with industrial processes and standards like IEC 62443.

Cisco Secure Equipment Access (SEA) provides controlled zero-trust secure remote access designed specifically for OT environments. It enables least-privileged access by allowing users to connect only to specific devices, using only approved protocols, and only during authorized times, aligned with the requester's responsibilities. SEA eliminates the need for traditional VPNs by embedding zero-trust network access (ZTNA) gateway functionality directly into Cisco industrial switches, simplifying deployment and scaling without requiring dedicated hardware or complex firewall configurations.

Cisco Identity Services Engine (ISE) is a comprehensive network security policy management platform that enables secure access control (via TrustSec in the network infrastructure) and visibility for users and devices across wired, wireless, VPN, and 5G networks.

Cisco Catalyst Center is a powerful network management system that leverages AI to connect, secure, and automate network operations. It simplifies the management of Cisco Catalyst network infrastructure, including IE switches, helping ensure a consistent user experience across wired and wireless networks.

Key design recommendations

To accelerate enterprise approval and reduce deployment risk, standardize the following as part of the Cisco Validated Design:

- **Dedicated segmentation for vision workloads** (camera/edge/management zones) with explicit policies for allowed traffic flows.
- **QoS policy baseline** that protects latency-sensitive industrial automation traffic while supporting sustained vision throughput.
- **PoE planning standards** (camera class, power budgets, redundancy considerations) to avoid rework during scale-out.
- **Time synchronization standards** (facility-wide) to ensure reliable correlation across cameras and operational systems.
- **Operational visibility and change detection** so OT teams can quickly identify new devices, unexpected communications, or policy drift.
- **Secure remote access standard** for internal support and vendors that avoids broad network exposure and eliminates “temporary” exceptions.

Summary

AI-driven machine vision is moving from experimentation to an enterprise capability—but scaling requires a secure, resilient, and repeatable deployment model. Invisible AI delivers operational intelligence from video to improve productivity, quality, and safety, while Cisco provides the industrial networking and security foundation that enables consistent performance, segmentation, visibility, and controlled access.

Together, Invisible AI and Cisco help manufacturers move faster—from pilot to multisite rollout—by reducing deployment friction, lowering operational risk, and meeting the expectations of enterprise IT and cybersecurity stakeholders.

C11-6063778-00 05/26