Cognex and Cisco Accelerating AI-Driven Machine Visioning

March 2026

This paper was authored by Bhavin Jethra (Cognex) and Paul Didier (Cisco) with contributions from Vivek Bhargava (Cisco)

Executive summary

AI-driven machine vision is increasingly used in manufacturing and logistics to improve throughput, quality, and operational visibility. Advances in sensors, edge computing, and AI have expanded the scope of machine vision applications beyond traditional inspection and code reading. Yet despite these advances, adoption at scale remains limited across many industrial environments.

The primary obstacles are no longer related to vision algorithms or model accuracy. Instead, organizations face challenges around **infrastructure readiness, deterministic system behavior, security, and operational trust**. AI-powered vision workloads introduce new requirements for time synchronization, bandwidth management, edge compute utilization, and protection of sensitive image data, capabilities that many production networks were not originally designed to support. These concerns often prevent organizations from moving beyond isolated pilots, even when the business value of machine vision is well understood.

This white paper presents a **joint Cisco and Cognex reference architecture** designed to address these challenges. By combining **Cognex's machine vision expertise** with **Cisco's industrial networking, security, and observability capabilities**, the architecture enables a structured, trust-first approach to deploying AI-driven machine vision in real production environments.

In addition to architectural guidance, the paper summarizes key validation observations that demonstrate how deterministic networking, traffic prioritization, secure segmentation, and centralized edge operations can reduce deployment risk and support scalability. The focus is on practical design principles that help organizations confidently progress from experimentation to production.

Business drivers for AI-driven machine vision applications

Manufacturing and logistics organizations face growing pressure to move faster, improve accuracy, and build resilience. Labor shortages, higher fulfillment expectations, and stricter compliance standards are driving a reassessment of work across production lines, sortation systems, and distribution networks.

Labor variability

Manual inspection, sorting, and verification tasks are difficult for staff to perform consistently and are prone to fatigue-related errors. Machine vision systems provide consistent, repeatable inspection and decision-making, enabling organizations to stabilize operations while reallocating human effort toward higher-value exception handling and supervision.

Throughput and accuracy requirements

In both manufacturing and logistics environments, small error rates can translate into significant downstream costs, including rework, mis-sorts, customer dissatisfaction, and regulatory exposure. AI-enhanced vision allows systems to handle greater variability in products, packaging, and environmental conditions without sacrificing speed or reliability.

Traceability and operational visibility

Customers, regulators, and internal stakeholders increasingly expect detailed records of what was processed, when decisions were made, and why. Machine vision systems can generate structured data and visual evidence that supports quality assurance, compliance audits, and continuous improvement initiatives.

Together, these drivers are pushing AI-driven machine vision from a niche optimization tool toward a foundational capability for modern production and logistics systems.

Target audience

This white paper is for teams deploying and operating AI-driven machine vision systems in manufacturing and logistics environments, where success depends on coordination across IT, OT, and operations.

IT and security teams evaluate how vision systems integrate with existing infrastructure, with a focus on deterministic networking, segmentation, visibility, and protection of sensitive data.

Operations leadership prioritizes reliability, uptime, and safe day-to-day operation, ensuring that new systems improve throughput and quality without disrupting production.

OT engineers and systems integrators design and implement real-time vision and automation systems, integrating cameras, edge compute, and control systems to achieve predictable behavior at scale.

By addressing the needs of these groups collectively, the guidance in this paper supports a unified, trust-first approach to machine vision adoption.

Critical challenges facing customers deploying AI-driven machine vision

AI/ML training and inferencing development

Moving from rules-based to AI- and ML-driven image processing is a significant shift. While the benefits can be substantial, training and deploying models require large volumes of representative data that is often difficult to extract from existing vision systems and networks. In many cases, the compute resources for training reside in the cloud or separate on-premises infrastructure (e.g., GPU servers or lab-based PCs), further complicating data movement and integration.

Infrastructure readiness

Many production networks were built for deterministic control traffic and basic automation—not high-bandwidth image data or AI edge workloads. As camera resolution increases and AI inference expands, concerns grow around latency, jitter, synchronization, and bandwidth contention. Without confidence that time-sensitive control traffic will remain protected, organizations hesitate to introduce vision into production environments.

Operational complexity

Operational complexity at the edge is a significant challenge. AI-driven machine vision depends on distributed industrial PCs, GPUs, and edge software that require ongoing updates, monitoring, and maintenance. Organizations fear creating fragile, hard-to-manage systems—especially across multiple lines or sites.

Protecting critical, sensitive Intellectual Property (IP) data

Security and data governance concerns further widen the gap. Vision systems process sensitive images that may expose proprietary products, processes, or customer information. IT and security teams require clear data protection, controlled communication paths, and assurance that new devices do not expand the attack surface. Uncertainty around data flows and access controls often delays or prevents deployment.

Lack of guidance

Finally, many organizations lack clear, safe-to-start deployment guidance. Without reference architectures or validated design patterns, teams are forced to make early architectural decisions with long-term consequences. The perceived irreversibility of these choices—combined with cross-functional ownership challenges between IT, OT, and operations—often leads to stalled initiatives.

Networking and security requirements to accelerate deployment

Preparing production infrastructure for AI-driven machine vision requires meeting the following key requirements.

Deterministic timing and synchronization are essential for machine vision systems that interact with high-speed production and material handling processes. Cameras, edge compute, and control systems must maintain accurate time alignment to ensure correct correlation between images, physical position, and control events. Time-sensitive traffic must be protected to prevent drift or misalignment under load.

Predictable network behavior under congestion is another critical requirement. Vision systems generate a mix of time-critical control traffic and high-bandwidth image data. Networks must prioritize synchronization and control traffic while efficiently transporting image streams without impacting system stability. Traffic classification, prioritization, and efficient packet handling are necessary to maintain consistent behavior as data volumes increase.

Efficient edge processing and data movement are required to manage the growing scale of visual data. AI inference should occur close to the source to reduce latency, limit bandwidth consumption, and keep sensitive image data under local control. Rather than transporting raw images indiscriminately, systems must support selective data extraction, filtering, and forwarding based on operational and business needs.

Operational visibility and manageability at the edge are key to sustaining deployments over time. Distributed vision devices, industrial PCs, and AI workloads must be centrally observable, configurable, and maintainable without introducing excessive operational overhead. Organizations require confidence that systems can be updated, monitored, and recovered in a controlled manner across multiple sites.

Security, segmentation, and data governance are prerequisites for deployment approval in modern industrial environments. Vision systems must integrate with existing IT and OT security frameworks, enforce explicit communication boundaries, and support identity-based access control. Sensitive image data and model behavior must remain protected throughout the system lifecycle.

Finally, **scalability and repeatability** are essential for long-term success. Architectures should support consistent deployment patterns across lines and facilities, enabling organizations to expand adoption without redesigning infrastructure or increasing operational risk. Clear separation of responsibilities between vision intelligence, edge execution, and infrastructure trust is fundamental to achieving this goal.

Cognex overview

Cognex is a global provider of machine vision systems used in manufacturing and logistics environments for inspection, identification, and process control. Its portfolio includes smart cameras, industrial vision systems, and barcode readers, that enable real-time visual interpretation and decision-making on the production floor.

Within the context of AI-driven machine vision, Cognex technologies provide the sensing, inference, and lifecycle governance capabilities required to deploy and scale vision applications across distributed industrial environments. Together, these capabilities form the Cognex vision-system layer of the joint reference architecture described in the following section.

Reference architecture for machine visioning in manufacturing and logistics production systems

Machine vision reference architecture

To address the requirements outlined above, this section presents a **joint reference architecture** that separates vision intelligence, operational execution, and trust in the infrastructure into clearly defined layers. This separation enables organizations to deploy AI-driven machine vision with predictable behavior, strong security, and scalable operations across manufacturing and logistics environments.

The architecture combines **Cognex's expertise in machine vision** with **Cisco's capabilities in industrial networking, security, and observability**. Each layer is independently governed, while working together as a unified system.

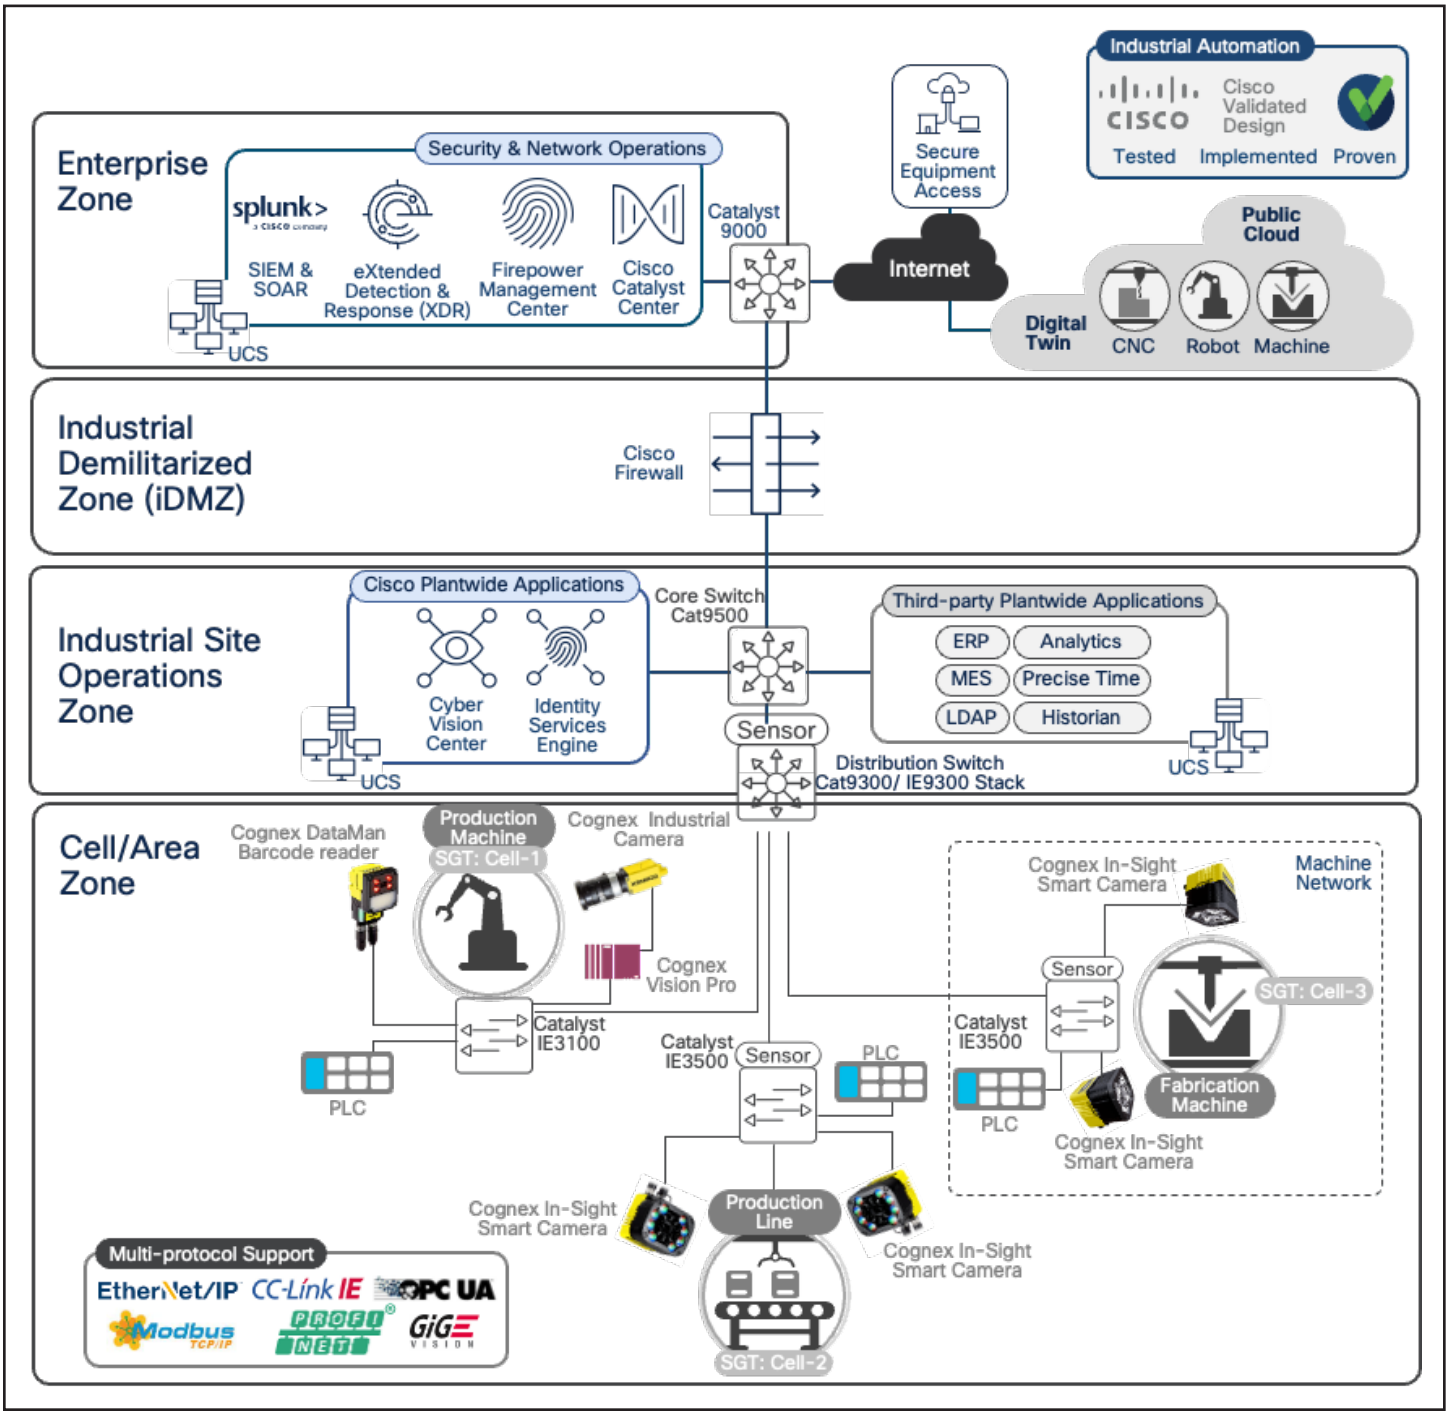

Figure 1. Machine Vision for Industrial Automation architecture map

Cognex components

The Cognex components in this architecture collectively provide vision sensing, AI-driven interpretation, edge-level execution, and centralized governance. Each component plays a distinct role while operating as part of an integrated machine vision system.

Figure 2. Cognex cameras and barcode reader

Cognex vision sensors and readers

Cognex industrial cameras and smart devices form the sensing layer of the system, capturing images and visual data directly from production and logistics environments. These include:

- **Cognex In-Sight smart cameras** for embedded vision inspection and guidance
- **Cognex DataMan barcode readers** for high-speed identification and traceability
- **Cognex industrial cameras with VisionPro** for advanced, PC-based vision applications

These devices perform real-time inspection, identification, and classification tasks at the edge, using both traditional vision techniques and AI-enhanced models where appropriate.

Together, these Cognex components ensure that vision intelligence is **executed locally, managed operationally**, and **governed centrally**, providing a robust foundation for AI-driven machine vision in production environments.

Cisco components

The key Cisco components include:

Cisco **Industrial Ethernet (IE)** switches: Available in DIN-rail, IP67-rated, and rack-mount form factors, 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, and offer high-wattage Power over Ethernet (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's Trust Anchor, Cisco Cyber Vision, Cisco TrustSec® segmentation, and Secure Equipment Access (SEA) for zero-trust remote management of connected assets. These switches offer high-capacity, low-latency Layer 2/3 switching, all managed by Cisco Catalyst Center.

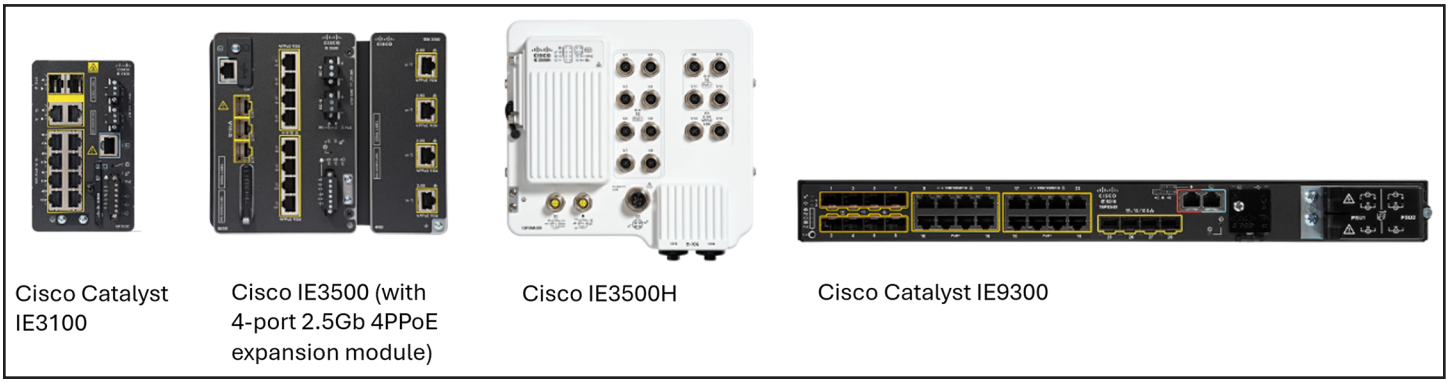

Figure 3. Cisco Industrial Ethernet switches

Table 1. Cisco Industrial Ethernet switches for machine vision cameras

IE switch family	Port speeds (count)			PoE			Jumbo frames	Cyber Vision, TrustSec, SEA
	10GE	2.5GE[1]	1GE[1]	4PPoE	PoE/PoE+	Budget		
IE3500	3	0/4	8/24	12		480 W	Yes	Yes
IE3500H	2	2	12	2	12	240 W	Yes	Yes
IE9300	4	8	16	8	16	720 W	Yes	Yes
IE3100	-	-	10	2	6	240 W	Yes	SEA only

[1] The number after the "/" indicates total ports available with module added to base switch.

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.

Validation highlights and design implications

Representative deployment scenarios were validated to assess the behavior of AI-driven machine vision workloads across industrial networking and edge environments. While specific measurements are omitted for confidentiality, the following observations highlight system-level behavior relevant to real-world deployments.

Deterministic timing and synchronization were confirmed as foundational to reliable machine vision operation. Accurate time alignment across cameras, edge systems, and control components enabled consistent correlation between image capture and physical processes, even as network complexity increased.

Network latency, jitter, and scalability remained stable as additional industrial switching elements were introduced. Properly designed multihop industrial networks did not degrade system behavior, demonstrating that machine vision deployments can scale across realistic facility topologies without compromising determinism.

Traffic management and bandwidth efficiency played a critical role in maintaining system stability. Separating time-sensitive control and synchronization traffic from high-bandwidth image streams, combined with qualified traffic prioritization, ensured predictable behavior during periods of elevated image throughput. Efficient packet handling further reduced processing overhead for high-resolution vision data.

Edge operations and visibility improved operational confidence by enabling centralized monitoring and management of distributed vision devices and workloads. Visibility into system health and data flows reduced the perceived risk associated with maintaining AI-driven machine vision at scale.

Security segmentation and asset visibility provided clear insight into connected devices and communication paths. Identity-based segmentation and controlled management access limited exposure while preserving required operational communication, reinforcing trust across IT and OT stakeholders.

Overall design implications

These validation observations demonstrate that AI-driven machine vision can behave as a predictable, secure, and scalable industrial subsystem when timing, traffic behavior, edge operations, and security are treated as first-class architectural concerns. Intentional system design enables organizations to move beyond experimentation and adopt machine vision with confidence in production environments.

Key network design considerations for AI-driven machine vision include:

- Use a **high-bandwidth** network infrastructure that supports multiple 10GE interfaces for uplinks and 1GE or 2.5GE interfaces for camera connectivity, such as the IE3500, IE3500H, and IE9300. For lower-speed camera connectivity, consider the IE3100.
- Use, where supported, PoE-capable switches that can power Cognex vision devices, reducing cable infrastructure and simplifying deployment.
- Use **Precision Time Protocol (PTP)** from the network to synchronize multicamera systems.
- Apply **Quality of Service (QoS)** policies that prioritize PTP, control traffic, and image transfer.
- Use network infrastructure that supports **jumbo frames**, such as the IE3500 for image transfer to lower latency and jitter.
- Place cameras and image processing GPU/IPC servers in the **same, dedicated Layer 2 domain/VLAN**.
- **Keep the switch hop count low in the vision path**.
- Use Cisco's TrustSec technology to create **zones and conduits (micro-segmentation)** to manage communication flows into and out of the image processing VLANs/domains.

For more details, refer to the Machine Vision in Industrial Automation Environments Design and Implementation Guide.

Summary

AI-driven machine vision has reached a level of technical maturity that enables meaningful improvements in throughput, quality, and operational visibility across manufacturing and logistics environments. However, broad adoption depends less on vision accuracy alone and more on the ability to deploy these systems **safely, predictably, and at scale** within real production constraints.

This white paper has shown that the primary barriers to adoption—deterministic behavior, infrastructure readiness, operational complexity, and security—can be addressed through a **system-level architectural approach**. By combining **Cognex's expertise in machine vision** with **Cisco's capabilities in industrial networking, security, and observability**, organizations can establish a trust-first foundation for AI-driven machine vision.

The joint reference architecture and validation highlights presented in this paper demonstrate how intentional design choices—such as deterministic timing, traffic prioritization, identity-based segmentation, and centralized edge operations—enable machine vision systems to behave as reliable industrial subsystems rather than best-effort IT workloads. These principles reduce deployment risk, clarify ownership across IT and OT teams, and support repeatable scaling across lines and facilities.

For organizations beginning their machine vision journey, this architecture provides a practical starting point. For those looking to expand existing deployments, it offers a structured path from isolated pilots to production-scale programs. By focusing on trust, predictability, and scalability from the outset, AI-driven machine vision can be confidently integrated as a core capability of modern industrial systems.

C11-5922295-00 03/26