



Functional Safety for Fenceless Robots

The trend towards fenceless robots in industry, and its implications for the future of functional safety in robotics design.



Contents

01	Executive Summary	3
02	The Industry 4.0 Paradox	4
03	The Automation Playbook	5
	Automation Model 1 - The Lights-Out Factory	5
	Automation Model 2 - The Cobot Compromise	6
	Automation Model 3 - Agile by Design, Constrained by Sensors	7
04	The Outside-In Trap	9
	The Occlusion Problem	9
	The Geometry Problem	10
	The Flexibility Illusion	10
	The Cost and Integration Problem	11
05	The Mantis Resolution	12
	The Inside-Out Architecture	12
	Solving Occlusion and the Geometry Problem	12
	Solving Integration Cost and the Flexibility Illusion	13
	The 4-Layer Safety System	14
	Safety Architecture and Certification	14
	Platform Independence: STAR Works on Any Robot Form Factor	15
	The Native Industry 4.0 Robot	15
06	Beyond Deterministic Safety	16
07	Conclusion	17
	References	18

01

Executive Summary

From fenced cells to embedded intelligence: a new architecture for human-robot collaboration, with uncompromising speed.

Manufacturing's labor crisis is structural: Deloitte and the Manufacturing Institute project 1.9 million unfilled U.S. manufacturing jobs by 2033, a pattern mirrored across Europe and Asia as Baby Boomer retirements accelerate and birth rates decline. Industry 4.0 offered an answer by delivering a dynamic, human-machine factory built for rapid changeover and mixed production. However, the physical safety infrastructure of the previous era continues to block its realization on the shop floor. Three automation models have emerged in sequence, each solving some problems while creating new ones: the caged lights-out cell, the collaborative robot, and the fenceless sensor-arbitrated cell. All three share a common architectural flaw: safety intelligence lives in the factory, not in the robot.

This white paper examines a novel architectural approach that embeds proximity sensing directly into the robot arm itself, making the safety function a property of the robot rather than of the cell it is deployed in. Decoupling the safety function from cell-level integration removes the validation burden that has, until now, made fenceless automation costly to deploy and difficult to redeploy. This directly results in a robot that arrives pre-certified, installs in days, and can follow the factory rather than dictating its layout, making it the first architecture genuinely compatible with the Industry 4.0 promise: **full industrial speed, true fenceless safety, and the agility to adapt to the factory** rather than demanding the factory adapt to it.

02

The Industry 4.0 Paradox

The Cage and the Promise: Why Industry 4.0 Is Stalling on the Factory Floor

The working-age population in every major manufacturing economy is shrinking. Baby Boomer retirements in the United States alone are removing 10,000 workers from the labor force every day. Deloitte and the Manufacturing Institute project that 3.8 million manufacturing positions will open between 2024 and 2033, with nearly half going unfilled due to a structural skills gap that no training program can close at scale. The labor force participation rate in the United States has fallen from 67.3% in 2000 to 62.5% in 2024, and the pipeline of workers willing to fill physically demanding, repetitive shop-floor roles is narrowing faster than the economy can replace them. Wage pressure compounds the problem: average manufacturing compensation has reached \$45.85 per hour including benefits, a wage premium that makes purely manual assembly lines economically untenable for all but the highest-value operations. Supply chain volatility, accelerated by geopolitical disruption and demand unpredictability, has made fixed, high-volume production configurations a strategic liability. A factory optimized to run one product at volume is a factory that cannot respond to a world that no longer orders in that pattern. Operational agility is no longer a competitive advantage. It is an existential requirement.

Industry 4.0 emerged as the answer to these converging pressures. Its central proposition was a redefinition of the factory floor: from a rigid, high-volume production line into a dynamic, human-machine environment built for rapid changeover, smaller batch sizes, and mixed production. The industrial robot, no longer isolated behind a cage, but integrated into a shared workspace alongside human colleagues, was the enabling technology at the center of this vision. In an Industry 4.0 factory, **humans and robots would each do what they do best, side by side, with safety managed through intelligence rather than physical separation.**

The physical steel cage and its modern equivalent, the external safety sensor perimeter, is the infrastructure that made industrial automation possible. However, it is now that same infrastructure that is blocking industrial automation's next evolution. Every approach to removing the physical fence has, so far, either introduced a productivity penalty severe enough to undermine the business case, or replaced the cage with a sensor architecture so complex and workcell-specific that the promised flexibility never materializes on the factory floor. Industry 4.0 is architecturally ready. The safety paradigm is not.



03

The Automation Playbook

Three Models of Industrial Automation: A Field Guide to What Exists Today

Automation Model 1 - The Lights-Out Factory

When Speed and Volume Made the Cage the Right Answer

The natural territory of Model 1 is **high-volume, low-mix production**: automotive body-in-white assembly, semiconductor fabrication, dedicated consumer goods lines. In these environments a high-speed robot operates inside a hard-guarded perimeter with interlocked access doors rated to PLd/Category 3 or higher. No human presence is required or expected during normal production. The machine runs continuously at its mechanical limits: speed, payload, and repeatability are entirely unconstrained by any need to share space. The physical cage was not an engineering compromise: it was the correct engineering answer to a specific problem. **When human presence is structurally unnecessary and volume justifies the fixed footprint, a cage delivers maximum throughput at minimum variable cost.** The incompatibility with modern manufacturing conditions is equally structural. A guarded cell **adds 4–8 m² of floor space** per robot, requires a fixed and validated configuration, and demands **Lockout/Tagout** procedures for every human access event, which turn a five-minute intervention into a thirty-minute procedure. When product variants multiply, batch sizes shrink, and the production plan changes weekly rather than annually, the economics and logistics of the caged cell collapse. It was built for stability. The market now demands the opposite.



Automation Model 2 - The Cobot Compromise

How the Industry's First Attempt to Remove the Cage Traded One Problem for Another

The collaborative robot or “cobot” was the industry’s first attempt to remove the physical cage from a broader class of applications. Its target was **small to mid-volume, mid-mix production**: electronics assembly, CNC machine tending, laboratory automation, and high-mix small-batch manufacturing where human judgement is regularly required in proximity to the robot. The enabling safety mechanisms are Safety-Rated Monitored Stop (SRMS) and Power and Force Limiting (PFL). SRMS halts the robot when a human enters a defined zone, and PFL allows physical contact by capping forces within the biomechanical thresholds originally specified in ISO/TS 15066 and now absorbed into ISO 10218-2:2025.

Together, they removed the need for a physical cage across a wide range of applications. Nevertheless, the **throughput bottleneck** is structural and unavoidable. To remain within safe contact force limits across any plausible collision geometry, a cobot carrying a meaningful payload must reduce its Tool Center Point (TCP) speed to approximately 0.25 m/s in collaborative mode, compared to 2 m/s or more for a caged industrial arm, a 4–5× throughput penalty. This is not a tunable parameter or a firmware issue, it is a direct consequence of the PFL safety philosophy itself. Any production line operating above roughly 30 parts per hour will find the economics unworkable. Cobots solved the cage problem and introduced a productivity cliff in its place.

The Evolving Standards Landscape

For most of the cobot era, practitioners working on collaborative robot applications had to navigate two documents simultaneously. ISO 10218 (Parts 1 and 2) governed the safety of the robot itself and its integration into a complete cell. ISO/TS 15066:2016, a technical specification rather than a full standard, provided the collaborative-specific detail that ISO 10218 lacked: the four collaborative modes (SRMS, hand guiding, PFL, and SSM), the biomechanical contact force limits for each region of the human body, and the risk assessment guidance for shared workspaces. ISO 10218-2:2025 resolves this. The updated standard formally absorbs the content of ISO/TS 15066, incorporating collaborative workspace requirements, biomechanical contact force limits, and risk assessment guidance for human-robot shared workspaces directly into a single normative document. For installations commissioned after 2025, ISO/TS 15066 is superseded.

The 2025 revision also introduces a terminology shift with real engineering implications. The terms “collaborative robot” and “collaborative operation” do not appear anywhere in the revised standard. “Collaborative application” replaces them throughout making clear that collaboration is a property of the application as deployed and validated, not of the robot as a product class. A robot is not collaborative. A validated application running on that robot, in that cell, for that specific task, is. This matters practically for how responsibility is shared. Under ISO 10218-2, the integrator bears responsibility for the risk assessment and validation of the complete robot application as installed. A robot certified under ISO 10218-1, as the MR-1 is, has its core safety functions pre-validated. But a non-compliant installation invalidates CE marking for the workcell regardless of the robot’s own certification. Pre-certification of the robot substantially simplifies the integrator’s task but it does not eliminate it.

Automation Model 3 - Agile by Design, Constrained by Sensors

The Fenceless High-Speed Cell and Its Hidden Costs

Model 3 addresses the use cases that neither caged cells nor cobots can serve: **high-mix, high-volume environments** such as e-commerce fulfilment, contract manufacturing, and dynamic kitting operations where product mix changes frequently and human operators regularly work in close proximity to the robot. The enabling safety mechanism is Speed and Separation Monitoring (SSM), defined in ISO/TS 15066 and incorporated into ISO 10218-2:2025. A surrounding layer of laser scanners, 3D time-of-flight cameras, and a safety PLC monitors human proximity in real time and modulates robot speed proportionally: full speed at safe separation distances, reduced speed as a human approaches, complete stop at the minimum separation zone.


Physical fences are replaced by a virtual, dynamically computed safety bubble. The promise is genuine: a full industrial arm, full industrial payload, full industrial speed, in a shared human space, without a cage. The hidden costs are equally genuine, and they arrive with the integration invoice.

Safety-Rated Monitored Stop (SRMS)

One of the four collaborative application modes defined in ISO/TS 15066 and incorporated into ISO 10218-2:2025. In SRMS, the robot is brought to a complete, monitored stop whenever a human enters a defined safety zone. The robot remains stationary until the zone has been vacated and the system has confirmed clearance, at which point normal operation resumes. SRMS does not limit robot speed or payload during normal operation, the safety function activates only on human presence. Appropriate where occasional human access is required (loading, inspection, and intervention) but continuous human-robot collaboration is not. The productivity cost is proportional to zone entry frequency: infrequent access has marginal impact, and frequent access makes the robot the constraint on line throughput.



HOW SPEED & SEPARATION MONITORING WORKS



“The promise is genuine. The hidden costs are equally genuine, and they arrive with the integration invoice.”

Power and Force Limiting (PFL)

One of the four collaborative application modes defined in ISO/TS 15066 and ISO 10218-2:2025. A robot operating under PFL continuously monitors and limits its own contact force and momentum so that, in the event of an unintended collision with a human, the resulting forces remain within the biomechanical injury thresholds defined in Annex A of ISO/TS 15066 (now Annex M of ISO 10218-2:2025). These limits specify maximum allowable contact force and pressure for each region of the human body, for both transient (impact) and quasi-static (clamping) contact scenarios. PFL allows genuine physical contact between human and robot without requiring a fence or zone detection, appropriate for tasks where direct physical interaction is part of the intended workflow. The structural limitation is speed: maintaining contact forces within safe limits across all plausible collision geometries requires TCP speeds of approximately 0.25 m/s or below in collaborative mode, compared to 2 m/s or more for caged industrial arms. This speed constraint is not a tunable parameter. It is a consequence of physics and the biomechanical limits themselves.



Speed and Separation Monitoring (SSM)

One of the four collaborative application modes defined in ISO/TS 15066 and ISO 10218-2:2025. In SSM, the robot's speed is modulated in real time as a function of the measured distance between the robot and any human in the workspace. A protective separation distance (PSD) is maintained at all times: the minimum gap required to guarantee the robot can be brought to a complete stop before contact occurs, accounting for robot stopping time, human approach speed, and measurement system response time, per ISO 13855. As a human approaches, the robot slows down. If the PSD is breached, the robot stops. SSM allows a standard industrial arm to operate at or near full speed in a shared human workspace with no physical cage, provided the protective separation distance is always maintained. The critical dependencies are perception quality (reliable, real-time, occlusion-free detection of human position), robot stopping time (which determines the required PSD and therefore the minimum practical working distance), and the accuracy and latency of the full sensing-to-stop signal chain.



04



The Outside-In Trap


Safety from the Outside In: Why External Safety Sensors Cannot Unlock True Fenceless Automation

The Occlusion Problem


External sensors operate on line of sight. The robot's own arm, the workpiece it is handling, the fixture it works on, and the tooling around it all create shadow zones that no floor-mounted or ceiling-mounted scanner can see. The arm itself occludes the scanner field as it rotates, creating blind spots that shift with every configuration and every motion. Overhead approaches (an operator leaning across a conveyor belt, or a worker reaching in from above) are invisible to ground-level 2D scanners.

The industry has been aware of this limitation for years, and the major robot manufacturers have tried to address it. ABB's SafeMove platform uses safety laser scanners mounted at the robot's periphery. The SWIFTI CRB 1300 integrates a safety laser scanner directly into the cell to enable fenceless operation alongside a high-speed industrial arm. Mitsubishi Electric and others have attached LiDAR units to the robot base, attempting to give the arm limited awareness of its immediate surroundings. These efforts reflect genuine industry consensus: fenceless, full-speed human-robot collaboration is the right direction. None of them resolves the fundamental geometry of the problem: **a sensor that does not move with the arm cannot see everything the arm can reach**. Addressing residual shadow zones typically requires 4–6 additional sensors at different heights and angles, each adding hardware cost, safety PLC logic complexity, and new potential for conflicting zone states that must be individually resolved and re-validated.


4-6 additional sensors typically needed to cover residual shadow zones



Overhead reach. An operator leaning in above the scan plane goes undetected.



Workpiece shadow. Stacked cases block the scanner's line of sight.




Arm occlusion. The robot's own arm blocks the scanner field.


The Geometry Problem

The principal argument for replacing a caged cell with a fenceless SSM system is floor space recovery: remove the fence, reclaim the perimeter. That recovery is real, but conditional on sensor placement that is rarely straightforward to achieve. The ISO 13855 protective separation distance formula computes the minimum safe gap between a detected human and the robot hazard zone as a function of human approach speed, robot stopping time, and sensing system latency. The closer the sensor is to the hazard, the tighter the required separation distance and the more floor space is genuinely recovered. Strategic sensor positioning is therefore not an aesthetic choice: it is the variable that determines how much of the promised floor space saving actually materialises. The difficulty is that optimal placement is constrained by the factory itself. Close overhead mounting can yield tight separation distances but requires a ceiling at the right height, which is absent in many industrial environments. Table or platform mounting introduces its own coverage gaps. Floor level mounting pushes sensors further from the hazard, increasing the required separation distance and eroding the saving the system was installed to deliver. A factory that replaces caged cells with fenceless sensor arbitrated cells to reclaim floor space may find, after accounting for what its facility actually permits, that the gap between the theoretical saving and the delivered saving is larger than anticipated.

“Sensor positioning is not an aesthetic choice: it determines how much of the promised floor space saving actually materialises.”



Caged cell. The fence sets the footprint.



Fenceless SSM cell. The footprint (dashed) depends on the ISO 13855 separation distance, and so on where the sensors can be mounted.

The Flexibility Illusion

The core commercial argument for fenceless external sensor systems is flexibility: move the robot, remap the zones, resume production. In practice this flexibility is heavily constrained. Moving a robot means moving its hazard zone, which requires recalculating ISO 13855 separation distances, repositioning sensors to restore coverage, and re-validating the entire safety certification. For high-mix environments where cell layout changes with product variants, the engineering overhead of maintaining a validated fenceless sensor system can negate the agility benefit entirely. Factories that anticipated weekly reconfiguration typically find they can manage it quarterly at best. The trajectory of Veo Robotics illustrates both the scale of the ambition and the difficulty of the approach. Veo pioneered production-ready,

vision-based fenceless safety with its FreeMove 3D safeguarding system, certified to ISO 13849 at PLd/Category 3, and backed by \$56 million from investors including Amazon, Google Ventures, and Yamaha Motor. FreeMove worked: it was a technically sophisticated, safety-certified system that enabled standard industrial arms to operate alongside humans without a physical fence. Its fundamental challenge was not the technology but the deployment model. Every new cell required sensor positioning, coverage validation, and safety function commissioning from scratch. Veo's assets were ultimately acquired by Symbotic in 2024, a company with the scale to absorb the integration overhead into its own fixed-installation infrastructure. Their technology found a home. The standalone, field-flexible product model did not.

The Cost and Integration Problem



Fenceless sensor-based cells are bespoke, multi-vendor integration projects: a safety PLC from one vendor, area scanners or 3D cameras from another, a robot controller from a third, a line PLC from a fourth. Each interface must be individually specified, wired, programmed, and validated. A full risk assessment and safety function specification under ISO 10218 and ISO 13849 is required for every cell from first principles; CE marking and technical file preparation must be completed from scratch. A firmware update to any single component can invalidate prior certification and restart the validation cycle. The numbers reflect this: a six-axis arm at roughly \$30,000 may require eight to ten 3D time-of-flight or LiDAR sensors at \$2,000–\$4,000 each just for spatial coverage, and the fully burdened cell cost — safety PLC, integration engineering, risk assessment, documentation, CE marking — routinely reaches 70% of total project budget. The robot arm, the component that does the actual work, is typically the least expensive item on the invoice. This is not a market inefficiency that competition will correct. It is a structural consequence of placing safety intelligence in the factory rather than in the robot.

Cost Item	Indicative Cost	% of Total Cell Budget
Robot arm (6-axis industrial)	~\$30,000	~15%
3D ToF / LiDAR sensors (8–10 units)	\$16,000–\$40,000	~15–20%
Safety PLC	\$3,000–\$8,000	~5%
Integration engineering	\$30,000–\$60,000	~25%
Risk assessment & documentation	\$15,000–\$25,000	~10%
CE marking & certification	\$10,000–\$20,000	~8%
Commissioning & validation	\$15,000–\$30,000	~12%
Total non-robot cost	~\$90,000–\$180,000	~70–85%

05

The Mantis Resolution

Safety from the Inside Out: How Embedded Intelligence Removes the Last Barrier to Industry 4.0


The Inside-Out Architecture

The limitations described above share a common root: safety intelligence is fixed in the factory while the hazard moves with the arm. The logical resolution is to move the safety function to where the hazard actually is. Rather than surrounding the robot with sensors, the sensing elements are embedded directly into the arm structure itself, distributed along the links and facing outward into the surrounding volume. Because the sensors are part of the arm, the protective sensing field envelops the robot continuously (from base to shoulder to end-effector) and moves with every joint position in every configuration. A human hand approaching from any direction, including from above or from an angle no floor-mounted scanner could cover, will enter the sensing field before it enters the danger zone. There are no geometric shadow zones, because there is no external sensor whose line of sight can be blocked. **The safety bubble is defined by the arm itself, moves with the arm, and covers the arm completely regardless of pose.**

Solving Occlusion and the Geometry Problem

Ceiling height, table configuration, and facility layout are irrelevant parameters, the sensors move with the arm. Because the sensing field is defined relative to the arm surface, the ISO 13855 separation distance is computed from a surface that is always co-located with the actual hazard, not from a fixed external reference point. The protective zone is tighter, more accurate, and does not generate the footprint overhead that external sensor systems impose.

“The safety bubble is defined by the arm itself, moves with the arm, and covers the arm completely regardless of pose.”




Solving Integration Cost and the Flexibility Illusion

The root cause of both problems is the same: when safety intelligence lives in the factory, every new deployment is a new safety engineering project. Mantis Robotics developed the MR-1 to resolve this at the architectural level. MR-1 is the first industrial robot arm to embed certified proximity sensing directly into its own structure, making the safety function a property of the robot itself rather than of the cell it is placed in. Because the safety function is certified at the robot level, not at the cell level, deploying the MR-1 does not require a new safety engineering project for each installation. Core safety functions are pre-validated. **The application risk assessment, which remains the integrator's responsibility under ISO 10218-2, becomes a verification exercise rather than a design exercise.** Redeployment by moving the robot to a new position, a new cell, or a new line, does not trigger a full re-validation cycle for the robot's safety functions. Only the application-specific parameters need to be reviewed. Deployment is measured in days, not months.

The 4-Layer Safety System

The MR-1 implements four independent safety layers, each with its own dedicated safety function. (i) Collaborative speed zone via SSM: the robot modulates TCP speed proportionally as a human approaches, maintaining full productivity at safe separation distances and reducing speed as the protective zone boundary is approached. (ii) Proximity stop via SSM: the robot comes to a complete, monitored stop when a human enters the minimum protective separation zone; the stop is triggered by the embedded STAR sensing system without reliance on any external sensor input. (iii) Active touch skin sensing via PFL: in tight-clearance applications where the arm passes close to fixtures or tooling and a worker's hand may be in the gap, an optional surface-mounted active-touch sensing layer across the arm exterior detects contact the instant it occurs and triggers an immediate stop. (iv) Joint-torque based PFL: joint-level torque monitoring provides hardware-level contact detection redundancy, independently of the external sensing layers. Each layer carries an independent safety function: the system does not depend on any single layer for its complete protective performance. The result is a defence-in-depth architecture: no single component failure results in an unmitigated hazard.



Safety Architecture and Certification

Each safety function in the MR-1 is designed to the required combination of Category 3 redundant hardware architecture, integrated diagnostic coverage across all relevant components, and MTTFd values appropriate to the required Performance Level. The complete product has been third-party certified under ISO 10218-1 and ISO 13849-1 and carries CE marking, confirming compliance with European machinery and safety directives. Because core safety functions

are certified at the robot level, the integrator's application risk assessment under ISO 10218-2 becomes a verification exercise to confirm that the application-specific parameters fall within the robot's certified operating envelope, rather than a multi-week safety engineering project starting from first principles. The division of responsibility is clear: Mantis certifies the robot and its safety functions and the integrator validates the complete application as installed.

Platform Independence: STAR Works on Any Robot Form Factor

A critical and underappreciated property of the STAR proximity sensing system is its architectural decoupling from the specific robot it is mounted on. The sensing elements are distributed across the arm's link surfaces and operate relative to those surfaces. The safety function is defined in terms of the proximity of external objects to the arm, regardless of the arm's absolute position in space or the configuration of its joints. This means the architecture is not fundamentally constrained to a single robot form factor. A six-axis vertical articulated arm, a SCARA, a delta, or a seven-axis redundant arm can in principle host the STAR sensing system, because the system measures proximity to the arm surface rather than proximity to a fixed external reference point. The safety bubble follows the geometry of whatever arm it is mounted on. This is a meaningful distinction from external sensor architectures, which must be redesigned and re-validated for every robot type, every cell layout, and every change in robot position. For system integrators working across multiple robot platforms and for manufacturers deploying different arm configurations across different lines, a safety architecture that travels with the arm substantially reduces validation overhead and simplifies the path to compliance across a mixed robot fleet.

The Native Industry 4.0 Robot

Every previous automation model asked factories to choose. The caged cell offered speed but no flexibility. The cobot offered flexibility but not speed. The fenceless sensor-arbitrated cell offered speed and a degree of flexibility but imposed an integration cost and re-validation burden that made that flexibility difficult to exercise in practice. **The MR-1 is the first robot that treats speed, safety, and agility as simultaneously non-negotiable.** There are no trade-offs to be managed. Because safety intelligence lives in the arm rather than the factory, the robot arrives pre-certified, deploys in days, and can be redeployed without restarting the entire safety validation process. It is not a cobot: it runs at industrial arm speeds with industrial arm payloads. It is not a caged cell: it operates without a physical fence in a shared human workspace. It is not a fenceless sensor system: it requires no external sensor infrastructure, making the cell inherently portable and reconfiguration-friendly. It is the first robot, architecturally designed for the Industry 4.0 factory, where the robot adapts to the factory rather than the factory adapting to the robot.

06

Beyond Deterministic Safety

AI Classification as an Advisory Layer

The safety architecture described in this paper is deterministic by design: the STAR proximity sensing system and its protective layers operate on fixed, certified logic whose behavior under any input is fully predictable and verifiable against ISO 13849. This is a requirement, not a limitation: functional safety certification currently demands it.

But deterministic safety systems are, by definition, context-blind. They know that an object is within the protective separation distance. However, they do not know whether that object is a torque wrench left on a fixture, a passing forklift, or a worker reaching in to clear a jam. The automotive industry has already navigated this boundary. In advanced driver assistance and autonomous vehicle systems, a certified deterministic safety layer (emergency braking, lane departure intervention, collision avoidance) operates independently of the perception stack. Sitting above it is a trained inference model that processes sensor data to build contextual awareness of the scene, distinguishing pedestrians from bollards, a merging vehicle from a stationary one, an open road from a construction zone. The classifier informs and anticipates, whereas the deterministic layer acts and stops. Neither replaces the other.

The same architecture is now a credible direction for human-robot collaboration. The proximity sensing data generated continuously by an embedded system like STAR is rich enough to train dedicated classifiers. These models learn to distinguish a worker's hand from a workpiece, a deliberate reach-in from an incidental approach, a routine cycle from an anomalous one. An inference model trained on this data could sit as an advisory layer above the deterministic safety core. It does not override it, but provides contextual signals that allow the robot to anticipate rather than merely react: slowing pre-emptively when a worker's posture suggests an imminent reach, holding a position when the scene is ambiguous, resuming confidently when context is clear. The regulatory and certification frameworks for AI-augmented safety systems in robotics do not yet exist in a form that would permit such a layer to carry a safety function in its own right. Standards bodies are aware of the gap, and ISO and IEC working groups are actively developing guidance on AI in functional safety contexts. The technology development, however, is already underway. The deterministic core certifies the present. The AI advisory layer is being built for what comes next.

07

Conclusion

The manufacturing labor shortage is structural, not cyclical. With 1.9 million U.S. manufacturing positions projected to go unfilled by 2033 and equivalent dynamics playing out across Europe and Asia, automation is no longer primarily a cost optimization decision. It is a matter of operational continuity. The question is not whether to automate, but which architecture can deliver automation at the speed, safety level, and flexibility that modern mixed-production environments actually require.

External sensor architectures for fenceless safety have validated the market direction and demonstrated genuine technical capability. They have not resolved the fundamental structural constraint: every deployment is a new safety engineering project, re-validation is mandatory after every reconfiguration, and the resulting cost structure and integration burden limit the practical flexibility of the cells they protect. Veo Robotics proved that vision-based fenceless safety could be certified and deployed at scale. The challenge was not the technology, it was the product model. Solving the cell is not the same as solving the problem.

Embedded safety intelligence on the arm is the architecture that resolves the constraint at its root. When the safety function moves with the robot rather than living in the factory, the robot arrives pre-certified, deploys without a bespoke sensor integration project, and can be redeployed without triggering a full re-validation cycle. The protective zone is complete, geometrically correct, and immune to occlusion by definition. The MR-1 is the first commercially available robot built on this architecture, certified to ISO 10218-1 and ISO 13849, carrying CE marking, and designed to deploy in days into the shared workspaces Industry 4.0 demands. The cage's era is not ending because safety has been abandoned. It is ending because safety has been solved at the right level of the architecture.

References

Bureau of Labor Statistics. (2025). Job Openings and Labor Turnover Survey (JOLTS). U.S. Department of Labor.

Cobot Safety DE. (2024). Biomechanische Grenzwerte bei MRK-Applikationen. Retrieved from <https://www.cobot-safety.de>

Coykendall, J., Hardin, K., Morehouse, J., Reyes, V., & Carrick, G. (2024). Taking charge: Manufacturers support growth with active workforce strategies. Deloitte Insights & The Manufacturing Institute.

EVS Robot. (2026). Cobot Safety Standards 2026: ISO 10218 vs ISO/TS 15066. Retrieved from <https://www.evsrobot.com>

IDEC Corporation. (2026). Behind the ISO 10218 series safety standards updates in 2025: understanding the real-world shifts in industrial robot technology that drove the revisions. Retrieved from <https://www.idec.com>

International Organization for Standardization. (2024). ISO 13855:2024: Safety of machinery — Positioning of safeguards with respect to the approach of the human body. ISO.


International Organization for Standardization. (2016). ISO/TS 15066:2016: Robots and robotic devices — Collaborative robots. ISO.

International Organization for Standardization. (2023). ISO 13849-1:2023: Safety of machinery — Safety-related parts of control systems — Part 1: General principles for design. ISO.

International Organization for Standardization. (2025). ISO 10218-2:2025: Robotics — Safety requirements — Part 2: Industrial robot applications and robot cells

Symbotic. (2024). Symbotic acquires assets of Veo Robotics [Press release]. Symbotic Inc.

Vu, C. (2020). Advanced vision and the future of robot safety in manufacturing. The Robot Report. Retrieved from <https://www.therobotreport.com>



FUNCTIONAL SAFETY FOR FENCELESS ROBOTS

Safety belongs inside the machine.

Robots already have brains and bodies. The next era of automation gives them reflexes, and leaves the fences behind.



Mantis Robotics, Inc. (HQ)
5673 W. Las Positas Blvd. Suite 215
Pleasanton, CA 94588 USA

+1-925-400-8083
info@mantisbots.com
www.mantis-robotics.com

Mantis Robotics BV
Gaston Geenslaan 6
3001 Leuven, Belgium

+32-16-29-83-80

Mantis Robotics Inc., Taiwan Branch
Rm. 8, 14F., No. 99, Sec. 1,
Xintai 5th Rd., Xizhi Dist.,
New Taipei City 221416, Taiwan (ROC)

+886-2-2697-5589