intenseye

SIF Prevention Playbook for the Containers, Packaging & Paper Industry

White paper

JUNE 2025

Table of contents

The SIF Dilemma in Packaging & Paper
Introduction
02

Global Packaging Company Scaling AI for SIF Prevention ROI
Case Study
16

From Safety-I and II to Safety-III
Embracing Real-Time Prevention
04

Analytics & Visualizing Safety Data
From Insights to Action
18

Major SIF Hazards in Packaging & Paper
Industry Focus
06

Implementing AI-Based SIF Prevention
Actionable Recommendations
22

Real-World Results in SIF Prevention
Case Study
14

A New Era of Proactive Safety Leadership
Conclusion
27

Introduction

The SIF Dilemma in Packaging & Paper

Serious Injuries and Fatalities (SIFs) remain a stubborn challenge in industrial sectors – and the containers, packaging, and paper industry is no exception. Over the past decade, overall injury rates have declined, yet the rate of fatal workplace incidents has not seen a similar drop.

In other words, minor incidents are down, but life-altering accidents continue to occur at distressing rates. **In 2023, over 5,000 U.S. workers lost their lives on the job (about 14 per day).** The paper and packaging sector has witnessed an alarming increase: in 2022 there were 13 fatalities in U.S. paper manufacturing (up from just 5 the year prior), with the majority linked to contact with dangerous equipment or exposure to harmful substances. This uptick is both ethically unacceptable and financially devastating – each fatality or disabling injury brings immense human cost, potential legal liabilities, downtime losses, and damage to a company's reputation.

Number of fatal work injuries

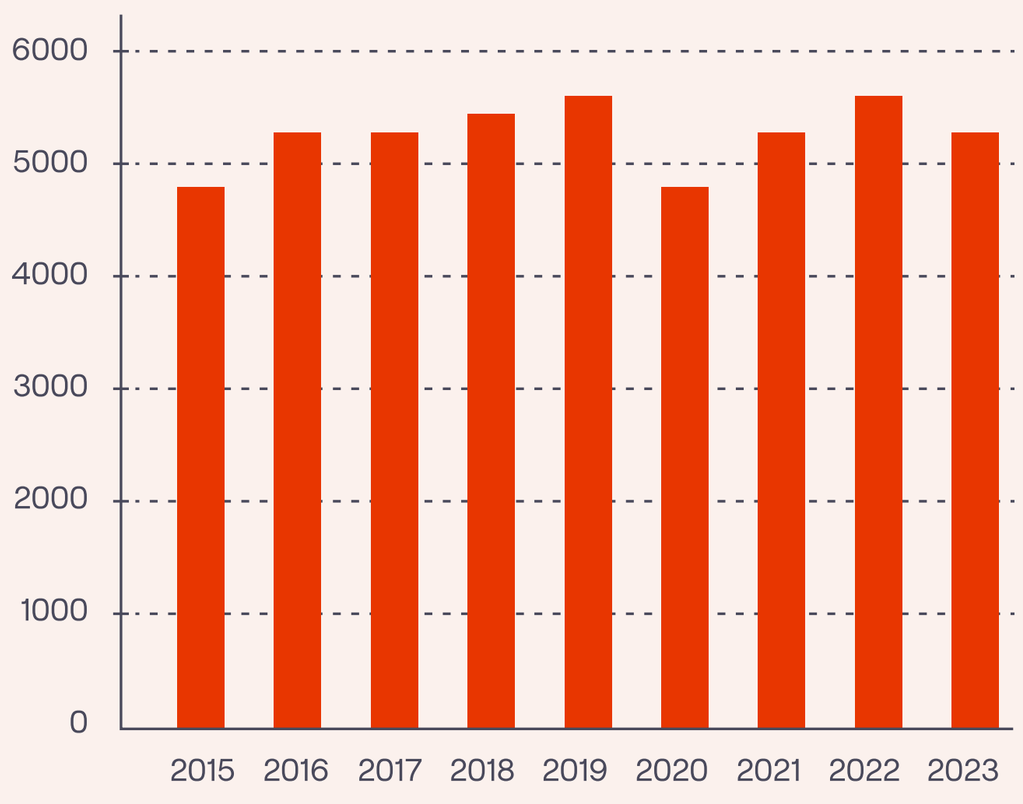

Source: U.S. Bureau of Labor Statistics

Why do SIFs persist even in well-established operations?

One reason is the gap between traditional safety metrics and SIF precursors. Many companies still rely on lagging indicators – OSHA recordables, after-the-fact investigations – which tell the story only after an injury has occurred. While Total Recordable Incident Rate (TRIR) has improved (U.S. TRIR fell from ~3.4 in 2012 to ~2.7 in 2022), the rate of fatalities plateaued around 3.0–3.2 per 100,000 workers in the same period. Within packaging and paper plants, teams often celebrate low injury rates, only to be blindsided by a serious accident that "came out of nowhere." Research shows the causes of severe incidents often differ from those of frequent minor injuries – meaning you can have a great overall safety record and still be at high risk for a SIF event.

Compounding the challenge, under-reporting of near misses and hazards is common: an estimated 79% of EHS leaders believe that unsafe situations are not reported consistently in their organizations. In busy corrugating plants or paper mills, employees may miss or even hide "small" incidents and close calls, not realizing these are flashing warning signs of a potential catastrophe.

The urgency for a new approach is underscored by sobering incidents in this industry. For example, in early 2024 a packaging worker in Washington state was crushed by a packing machine that unexpectedly energized – investigators found that permanent guards had been removed years prior and no effective lockout was in place.

In another case, a Houston employee was performing maintenance inside a cardboard baler when a co-worker unknowingly turned the machine on; the worker was tragically crushed to death on the spot. These aren't isolated "freak accidents" – they are the predictable outcome of control failures.

In the Houston baler incident, a basic lockout procedure could have saved a life. In a Wisconsin case, a 23-year-old lost three fingertips clearing a jam on a corrugator machine that cycled automatically; OSHA determined proper training and lockout devices would have prevented the amputation. Often there are red flags before disaster strikes: in one plastic packaging plant, two similar jam-up incidents occurred in the days before a fatal crushing – workers barely escaped injury in those near-misses.

Each of these tragedies carries a clear message: traditional safety programs and sporadic audits are not catching these SIF precursors in time. We need to move **from reactive hindsight to proactive, real-time foresight.**

The business case for action is equally compelling. SIF incidents carry huge direct costs (medical, compensation, regulatory fines) and indirect costs (production interruptions, retraining, morale damage). Industry leaders like the American Forest & Paper Association have set ambitious goals – for instance, AF&PA aspires to zero injuries by 2030 – recognizing that beyond moral duty, a strong safety record is a competitive differentiator. Companies with proactive SIF prevention see benefits in productivity and employee engagement: fewer emergencies mean more reliable operations, and a reputation for safety helps attract and retain a quality workforce. In sum, the status quo is unacceptable both ethically and operationally. The packaging and paper sector must confront its SIF problem head-on, leveraging new strategies and technologies to ensure every worker goes home safely.

This playbook introduces one such paradigm shift – **moving to "Safety-III" real-time safety management** – and lays out a practical roadmap for executives and safety professionals to drive serious injury and fatality prevention in their organizations.

From Safety-I and II to Safety-III

Embracing Real-Time Prevention

Breaking through the SIF plateau requires evolving how we think about safety. We've progressed from **Safety-I** (the old-school approach of reacting to incidents and measuring safety by the absence of accidents) to **Safety-II** (the more modern approach of being proactive and measuring safety by the presence of resilience and safe practices). Safety-II, championed by experts like Erik Hollnagel, encouraged learning from what goes right – near misses and normal work – not just what goes wrong. This shift led to improvements like empowering workers to speak up and designing systems with more tolerance for error. Yet, despite these advances, fatal and life-threatening incidents have not fallen in parallel; in some industries, serious incidents even increased while minor injuries declined. Traditional methods and even enhanced reporting under Safety-II have proven insufficient to eliminate SIF events. As one safety study noted, the next big step is achieving "greater clarity about how to identify and measure hazards in real time to intervene before incidents occur". In other words, we need to catch the precursors to a catastrophe as they are unfolding, not just analyze them after the fact.

Enter **Safety-III**, a paradigm that builds on Safety-II by adding real-time, tech-enabled risk management to the toolkit. Safety-III (a term our industry is using to describe this emerging approach) leverages modern technology – artificial intelligence (AI), computer vision, IoT sensors, advanced analytics – to monitor the workplace continuously and intervene before an accident happens. Think of it as shifting from looking in the rear-view mirror to having an always-on GPS and collision avoidance system for safety. In a Safety-III model, the focus is on **leading indicators and even live indicators of risk**.

Instead of waiting for an injury report or a near-miss card, the system watches for hazard signals in real time: a worker not wearing required PPE, a forklift and pedestrian getting too close, an unsafe entry into a machine zone – and it triggers immediate alerts or corrective actions. The goal is an active loop of detect → alert → correct, continuously running in the background of operations.

How does this look in practice? Imagine a network of smart cameras and sensors acting as a 24/7 "second set of eyes" in the plant. Computer vision AI analyzes live video feeds to recognize if **machine guards are left open**, if a person enters a restricted area, **if someone slips or a spill occurs, or if a forklift** is approaching a pedestrian too fast. The moment a dangerous situation is identified, the system notifies the relevant people – or even triggers an automated response like machine shutdown – so that the hazard can be addressed immediately.

For example, if an employee climbs onto a machine without locking it out, an alarm can sound before they start servicing it; if a walkway is blocked or a chemical leak is detected, alerts can prompt an instant response rather than waiting for someone to notice. Essentially, Safety-III augments human vigilance with AI-driven monitoring. It's about moving from lagging metrics (injury rates) to leading metrics (near-misses, unsafe conditions) – and now to **real-time metrics** (current hazards and compliance in the moment) to drive decisions.

It's important to note that Safety-III doesn't replace Safety-I or II principles – it enhances them. We still need a strong safety culture, training, and incident response (foundations of Safety-II). But Safety-III adds a continuous layer of prevention that addresses the critical gap: those infrequent but catastrophic scenarios that slip through our existing nets. By instrumenting the workplace and using AI and analytics to actively manage risk, organizations can achieve what was previously elusive – **preventing the "one-off" freak accidents** that devastate lives. In summary, Safety-III is about being proactive at the speed of operations. It represents a shift from retrospective analysis to real-time intervention. Companies adopting this model are seeing that it not only averts harm but also fosters a stronger safety mindset – workers become more aware of risks when they receive immediate feedback, and leadership gains visibility into problems that were previously invisible. The following sections will explore how Safety-III applies to the specific SIF hazards of the packaging and paper industry, and how real-time prevention is being implemented on the factory floor.

Industry Focus

Major SIF Hazards in Packaging & Paper

Real-time SIF prevention must be tailored to the actual high-severity hazards of the **containers, packaging, and paper sector**. From corrugating plants and printing presses to pulp and paper mills, certain scenarios consistently account for the most serious injuries and fatalities. Below we highlight 5 "SIF hotspots" in this industry – what can go wrong, and how a Safety-III approach can address each:

1 Machine Guarding and LOTO Failures

Packaging and paper operations use powerful machinery – corrugators, die cutters, printing presses, slitters, rewinders, palletizers, balers – that can cause amputations or deaths if proper safeguards are absent. A huge share of SIF incidents in this sector involve workers getting **entangled, crushed, or struck by moving machine parts**. Common causes include bypassed interlocks, removed physical guards, and failures in Lockout/Tagout (LOTO) during cleaning or maintenance. For instance, Washington state authorities recently fined a paper mill $650,000 after a 32-year-old worker was crushed by a packing machine – investigators found safety guards had been removed and the machine was not adequately prevented from energizing during servicing. In another case, a technician was fatally pulled into a conveyor baler because a co-worker re-started the machine while he was inside – a direct result of not locking out the equipment. **These incidents show how small lapses (a missing guard, one step skipped) can have fatal outcomes.**

A Safety-III approach significantly reduces these risks. Computer vision can continuously ensure that machine guards and access doors are in place and closed whenever equipment is running.

Virtual “danger zone” monitoring can be set up around hazardous machinery – if a person’s hand or body enters a defined zone (for example, reaching into a cutter or nip point), an alert is instantly raised or the machine can even auto-stop. Real-time systems also enforce LOTO compliance: AI cameras can detect if someone is working on a machine without the power being isolated (e.g. seeing movement in a zone when the machine shows running status) and send an immediate warning. They can verify the use of required PPE during maintenance – for example, confirming that **face shields, cut-resistant gloves, or arc-flash gear** are worn when working on certain equipment. Notably, these platforms can even track procedural requirements like minimum team size for high-risk tasks: if a job requires two authorized people (a “buddy system”) for safety, the AI will flag if someone is working alone in that area.

All these controls create a constant safety net, so even if a worker is tempted to take a shortcut, the system will call it out. The result is a dramatic drop in machine-related SIF precursors. In fact, plants that implemented AI “line of fire” detection around dangerous machines have reported near-elimination of close calls – for example, one facility set up alerts for anyone stepping under a suspended load or into a robot’s swing radius, and **immediately corrected unsafe entries**, virtually ending those high-potential incidents.

> By watching what humans can miss, real-time monitoring prevents the momentary lapses or miscommunications that underlie so many machine accidents.

2 Vehicle and Forklift Struck-by Incidents

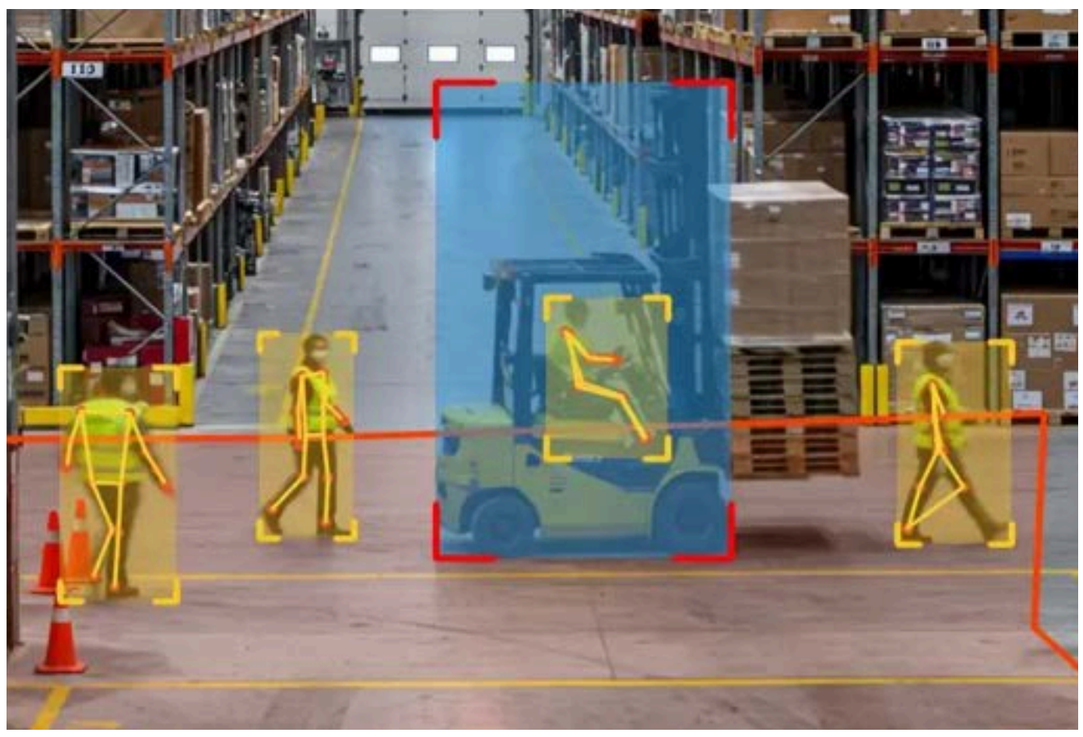

Busy packaging plants and warehouses often have internal traffic – forklifts shuttling palletized materials, roll clamps carrying giant paper rolls, yard trucks at loading docks, even Automated Guided Vehicles (AGVs) in modern facilities. Tragically, “struck-by” accidents (vehicles hitting a person) are a leading cause of workplace fatalities across manufacturing. In packaging and paper, we see scenarios like forklift vs. pedestrian collisions in warehouses, clamp trucks accidentally dropping huge paper rolls, or employees hit by semi-trailers during yard maneuvers.

In 2022, transportation incidents were a significant portion of manufacturing fatalities (for example, the plastics/rubber products segment saw six transportation-related deaths). One grim example: at a foam packaging plant in Oregon, an unsecured 2,200-pound bag of materials fell off a forklift and crushed a worker to death – OSHA found the load wasn’t properly secured and the operator hadn’t been adequately trained.

> Another incident involved a warehouse employee pinned between a giant paper roll and a forklift mast when a driver didn't see her. To address these hazards, companies are establishing **"virtual barriers" and AI-powered traffic management**.

Computer vision AI can monitor all forklift and vehicle movements in real time, issuing alerts when a pedestrian comes within an unsafe distance of a moving forklift, or if a forklift approaches a restricted zone too fast.

For example, **Intenseye's platform** uses stereo cameras and LiDAR to enforce defined safety distances – if a worker is, say, within 10 feet of a forklift that's not stopped, an audible alarm can warn both parties. The system can also detect violations like forklifts speeding, not using horn at intersections, or going the wrong direction, and immediately flag those to supervisors. This real-time coaching has an immediate impact:

Companies have seen metrics like a **20% drop in forklift speeding incidents within the first week** of deploying AI monitoring.

Another powerful tool is **heatmaps of near-misses**. The AI logs every close call (e.g. a too-close encounter or sudden stop) and aggregates data to show where in the facility these happen most frequently. One global packaging manufacturer discovered via heatmaps that a particular blind corner near the palletizer had an outsized share of close calls. In response, they reconfigured that area's layout and added AI alerting; the result was **zero incidents at that hotspot thereafter** and a measurable reduction in overall "risk exposure time" between pedestrians and forklifts.

This data-driven approach turns what was once a black box of forklift activity into a controlled, transparent process. Forklift operators become more accountable (knowing unsafe maneuvers are immediately flagged), and pedestrians gain an automatic guardian looking out for them. The bottom line: real-time monitoring of vehicle interactions can virtually eliminate the most catastrophic struck-by scenarios, while also improving material flow (fewer accidents mean less downtime).

3 Overhead and Dropped Object Hazards

In paper and packaging facilities, heavy materials are often stored or moved at height – think of massive paper rolls hoisted by cranes, pallets of product stacked high, or big bulk bags and racks in warehouses. This creates a risk of objects falling from overhead and causing serious injuries. Past incidents include workers crushed by paper rolls that fell from storage or clamps, and tools or parts dropped from an overhead conveyor or mezzanine. Even a lightweight object can be lethal if dropped from height (a wrench falling 30 feet can hit with thousands of pounds of force).

One case already mentioned was the 2,200-lb bag that slid off a forklift's raised forks and struck a worker. Another involved **a paper mill worker in Philadelphia killed when a huge roll of paper (over 1 ton) fell off a stack** – a catastrophic event. Contributing factors often include improper stacking, failure to secure loads, or workers standing under suspended loads (which should never happen). To counter these dangers, AI systems enforce "line of fire" rules in vertical space.

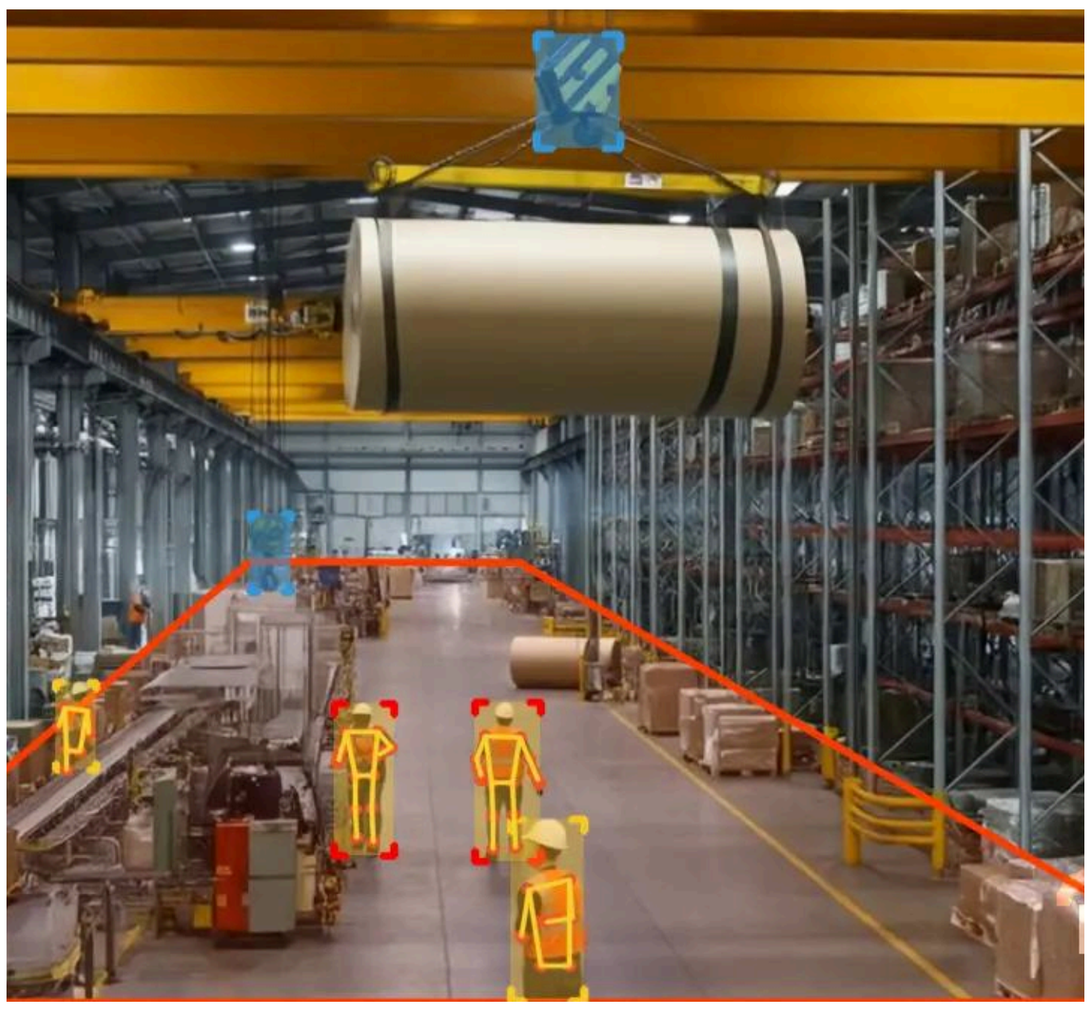

For example, **overhead crane zones** can be geofenced so that if a person walks beneath a suspended roll or bundle, an alert is triggered to both the pedestrian and the crane operator (or the crane is halted). AI can also monitor storage racks for safe stacking – identifying if a pallet or roll is protruding dangerously or if stored material looks unstable.

In automated storage and retrieval systems, vision algorithms verify that robotic handling of rolls/boxes is secure (no mis-gripped item). Additionally, **unauthorized zone entry detection** can prevent incidents: if an area under maintenance (e.g. a hoist area or high bay) is closed off, cameras will alert if someone enters without authorization. Some plants use **wearable tags in combination with AI cameras**: tags on hardhats can be sensed to shut down a hoist if a tagged worker is in a danger zone. The AI effectively creates a constant watch that no person is in harm's way when heavy objects are moving overhead. Furthermore, the system can detect "unsafe behaviors" like an employee standing on a forklift pallet to lift them (an extremely dangerous makeshift lift) – by training the AI on what normal vs. abnormal work at height looks like, such acts can be caught and stopped immediately. The benefit is twofold: protecting workers from lethal gravity-related accidents and reinforcing a culture that unsafe shortcuts (like walking under a lift or not securing a load) are never acceptable.

These measures help ensure that "what goes up, doesn't come down" unexpectedly on your workforce.

4 Fire, Explosion, and Hazardous Exposure

Another major SIF category in this industry involves **fires, chemical releases, or explosions** – low-frequency but high-consequence events. Packaging and paper facilities may not use as many toxic chemicals as, say, a chemical plant, but there are still significant risks: flammable liquids (inks, solvents, adhesives), combustible dust (paper dust or cornstarch used in corrugation can ignite), and common industrial chemicals (acids for cleaning, ammonia in refrigeration systems at some pulp/paper operations). Hot work (welding or cutting) during maintenance is a well-known ignition source.

The industry has seen devastating examples. In September 2020, at an Evergreen Packaging mill in North Carolina, a welding repair ignited vapors inside a process vessel – a flash fire killed two contractors and injured others. In another incident, a printing plant technician was cleaning a mixer with a solvent near a running UV curing unit; fumes ignited and caused an explosion, killing one worker and badly burning another. These scenarios underscore how a momentary lapse (failing to ventilate, or performing hot work without precautions) can turn into a fatal disaster. **Safety-III solutions** tackle these hazards on multiple fronts.

Firstly, environmental sensors (for gas, heat, smoke) can integrate with the AI platform: if a flammable vapor or toxic gas starts accumulating, the system raises the alarm before it reaches critical levels. For instance, fixed ammonia detectors in a refrigeration room can be tied to an AI dashboard that not only triggers evacuation alarms at the first sign of a leak, but also cross-checks camera feeds to ensure everyone evacuates and no one re-enters without clearance.

Secondly, computer vision monitors work activities for compliance with fire safety protocols. If someone is performing hot work (welding, grinding) outside of a designated, pre-cleared time and place, the AI can detect the welding arc or sparks and immediately alert safety managers. It also checks that during any authorized hot work, the required precautions are visible – e.g. is there a fire watch person present, is the area cordoned, are fire extinguishers nearby, is a welding curtain in place?

Intenseye's AI, for example, is being trained to recognize when an operator strikes an arc or uses a cutting torch, and if a fire blanket or curtain is missing, **it will flag it in real time.** Similarly, for hazardous chemical handling, the AI ensures proper PPE: if a worker is pouring a chemical and not wearing chemical-resistant gloves, apron, or face shield, an alert is generated so they can be stopped and corrected.

Another innovative application is **time-based monitoring** in dangerous areas – the system can track how long individuals spend in a confined or exposure-prone area. For example, if a technician enters a confined space (like a pulp digester or an ink mixing vat room), a timer starts; if they exceed the safe time limit, a warning is sent to check on them (to prevent overexposure or entrapment). By catching deviations (like an unexpected hot work, a missing gas detector, a worker staying too long in a solvent room), real-time systems maintain strict control over fire and chemical risks.

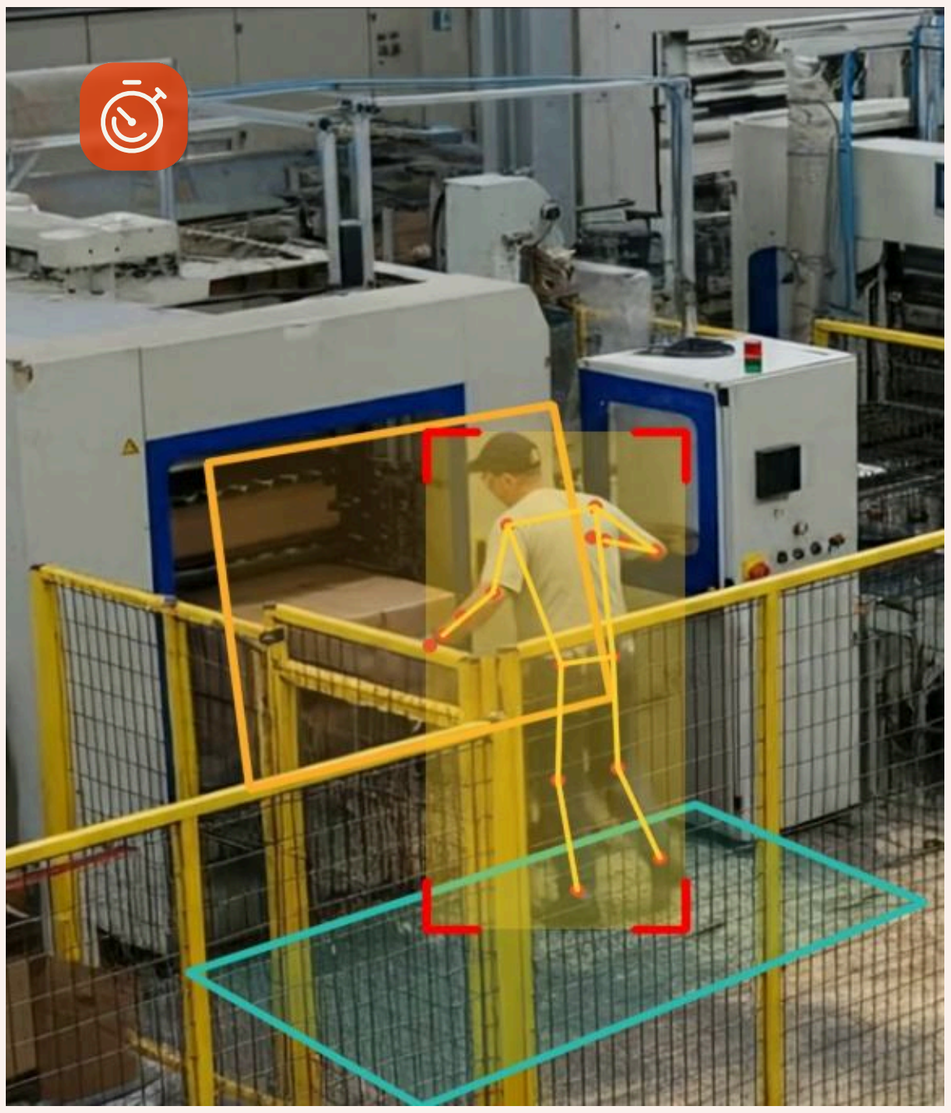

The outcome is fewer surprises: small fires or leaks are caught before they escalate, and employees are kept out of harm's way through automatic vigilance. It's like having a continuous Process Safety Management enforcer on duty – essential since even a single spark or whiff of toxic gas, if unnoticed, can have irreversible consequences.

5 Slips, Trips, and Falls

While often viewed as a common safety issue, slip/trip/fall hazards can lead to severe injuries and occasional fatalities, especially falls from height. In packaging and paper environments, there are unique contributors to these hazards. Paper mills, for instance, have wet processes – water or pulp slurry can make floors slippery.

Converting plants generate **paper dust and trim pieces** that scatter on floors, creating trip hazards if housekeeping doesn't keep up.

There are also many elevated work areas: catwalks on large machines, ladder accesses to mezzanines and silos, platforms around pulping or printing equipment. A fall from any of these (even a ~6–8 foot platform) can be deadly or cause permanent injury.

Nationally, **falls remain a top cause of occupational death (nearly 700 fall-related fatalities in the U.S. in 2022)**, and in food manufacturing (a comparable sector) falls accounted for 12 out of 75 fatalities in one year.

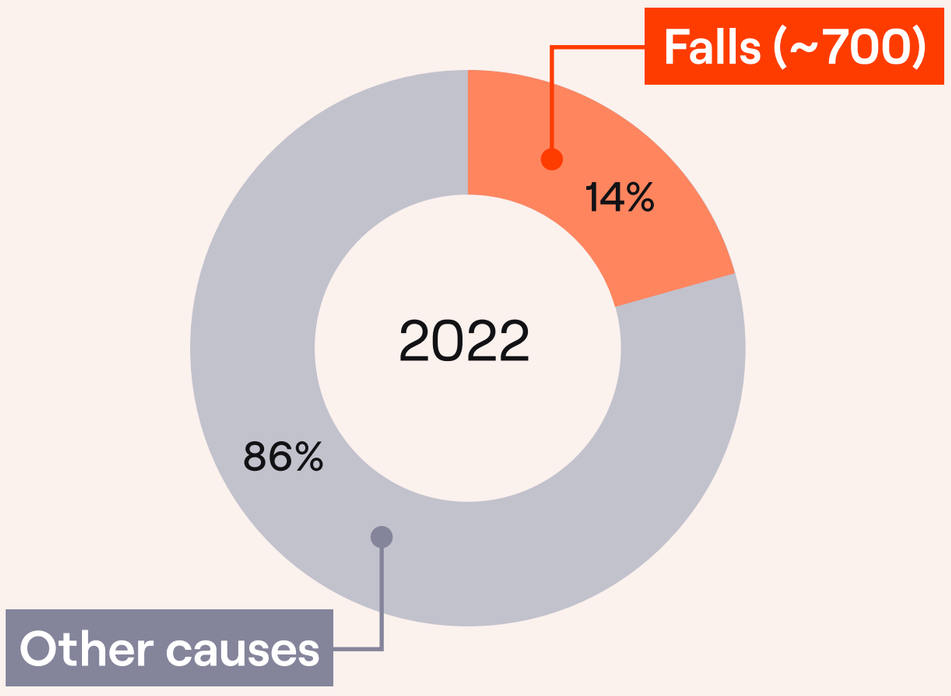

Occupational Fatalities (~5000 total)

In our sector, falls may be less frequent but have occurred – e.g. workers have fallen into paper pulpers or off warehouse racks while retrieving materials. Even same-level slips can cause major injuries like broken hips or head trauma (a slip on spilled coating chemical could lead to a concussion or worse). The Safety-III approach treats every uncontrolled slip or trip as a potential SIF precursor, not a minor issue. AI cameras can be taught to recognize **spills, leaks, and wet floor conditions** – for example, detecting a pooling liquid or powders on the ground – and immediately alert cleaning crews. One packaging facility using AI to flag spills saw a dramatic improvement: forklift oil leaks and hydraulic fluid drips (previously sometimes unnoticed for hours) were cleaned within minutes of occurring, resulting in a sharp drop in slip incidents.

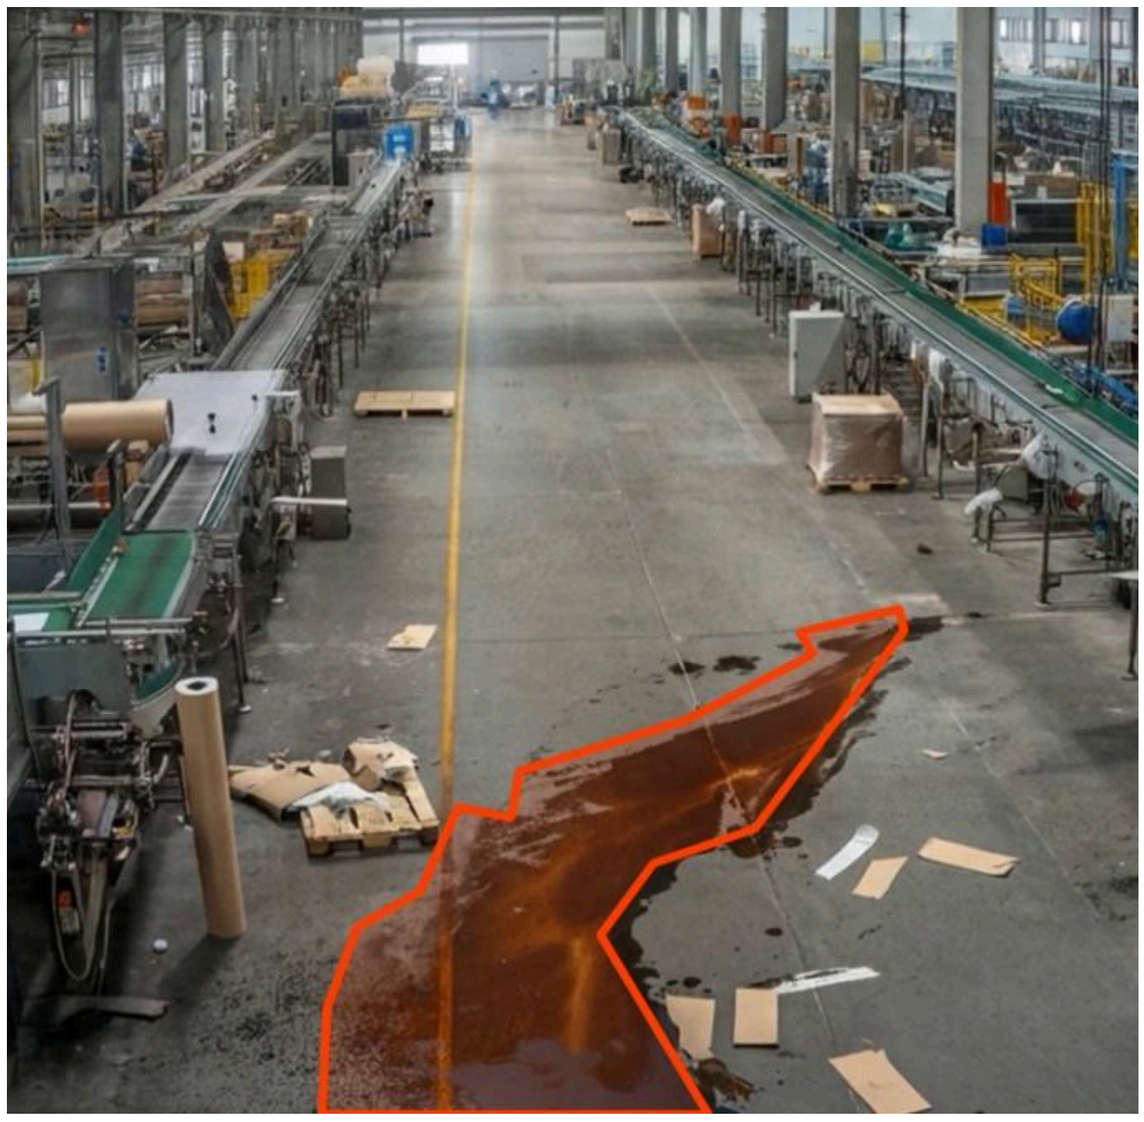

Intenseye's AI is also capable of detecting **when a person slips, trips, or falls** by analyzing unusual body postures/movements on camera. If a worker falls to the ground, the system can send an instant emergency notification – speeding up response when seconds count (especially important if someone falls in an isolated area). Moreover, to prevent **falls from height**, the technology can monitor if portable ladders are being used safely (e.g. angle, top contact) and whether fall protection is worn where required. For instance, if someone climbs on top of a machine or trailer without a harness where one is mandated, the AI will catch it and sound an alarm. By addressing the "early-warning" signs (spills, unsafe ladder use, unsecured work at height) in real time, companies can mitigate the most common injury in the workplace before it causes a broken bone – or a fatal fall. In fact, one Intenseye customer in food manufacturing credited real-time spill detection with creating a significantly safer floor environment, likely **avoiding major injuries** that would have otherwise occurred. The takeaway: even though slips and minor falls might not grab headlines, preventing them is crucial to reducing the overall serious injury burden, and AI provides tools to do it proactively.

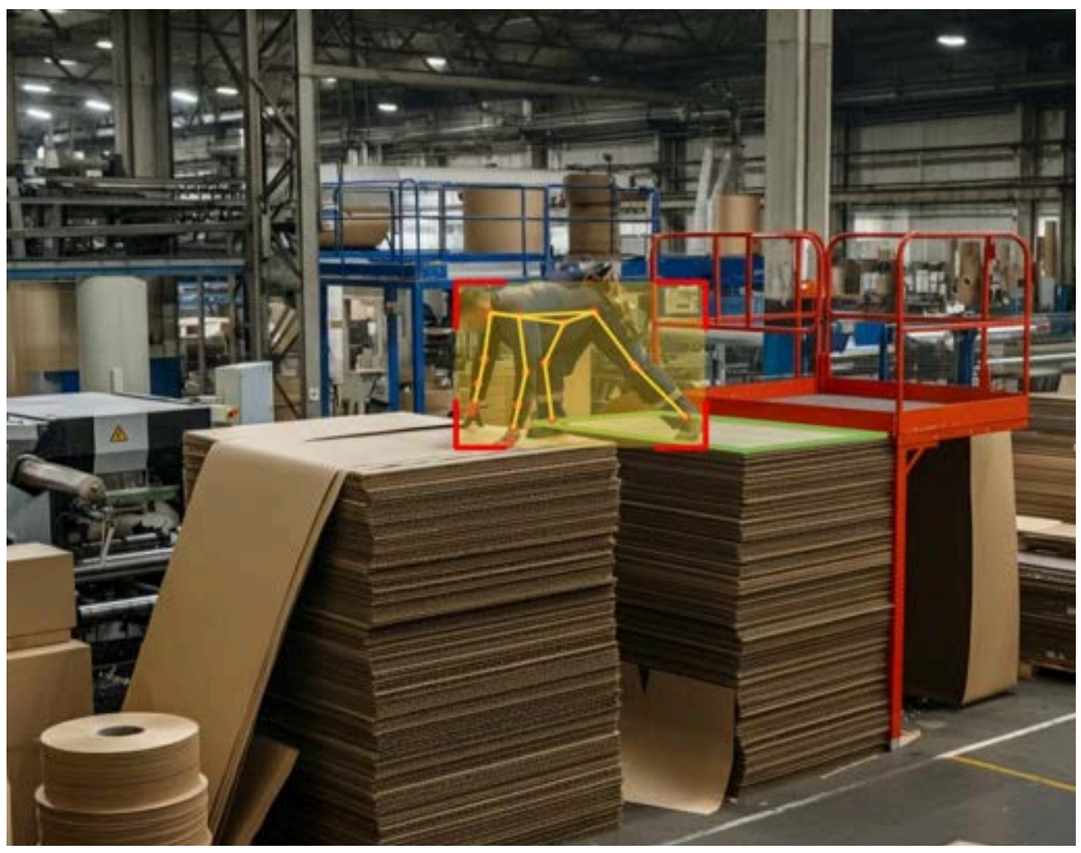

Each of the above hazard categories represents a significant SIF exposure in the packaging & paper industry, but **all are addressable** with the right mix of engineering controls, training, and real-time monitoring.

> A key principle of SIF programs is recognizing that **not all safety incidents are equal** – a simple rule violation in one context might be a nuisance, while in another it's a potential fatality.

For example, an employee not wearing earplugs is a compliance issue; an employee not locking out a machine is a life-threatening issue. The challenge for safety leaders is to ensure the truly critical situations are never lost in a sea of minor infractions. This is where AI-based systems excel: they can filter and prioritize a flood of observations by risk severity. Modern vision platforms categorize alerts as low, medium, or high based on the likelihood of serious injury – for instance, a worker missing a high-visibility vest might get a "low" alert, whereas a worker in a forklift path or a missing guard on a running machine triggers a "high" alert requiring immediate intervention. By automatically focusing attention on the SIF precursors, real-time systems help **EHS teams work smarter, not just harder.** Instead of spending hours on generic inspections or paperwork, teams can zero in on the 1% of situations that pose 90% of the catastrophic risk. The next section will show how this comes to life through case studies – companies in the packaging sector that transformed their safety outcomes using AI-driven, Safety-III methods.

Real-World Results in SIF Prevention

Case Study

Immediate Impact with AI-Powered Safety

ABOUT

Amcor Rigid Plastics, one of the world's largest packaging companies, has been a front-runner in adopting real-time safety management.

INDUSTRY

Packaging Industry

COMPANY SIZE

+41,000 employees globally

In 2022, Amcor deployed an AI-driven EHS platform across multiple production sites to target SIF precursors that were historically hard to monitor. The results were both immediate and significant.

Within the first week of use, Amcor recorded a **20% drop in forklift speeding incidents**, as the real-time alerts prompted drivers to improve behavior.

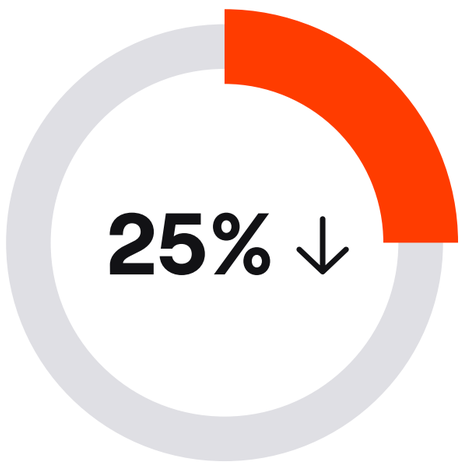

Near-misses that previously went unnoticed were now visible – and correctable. Over the course of a year, Amcor saw a **25% reduction in Total Recordable Incident Rate (TRIR)** at the equipped sites, compared to baseline.

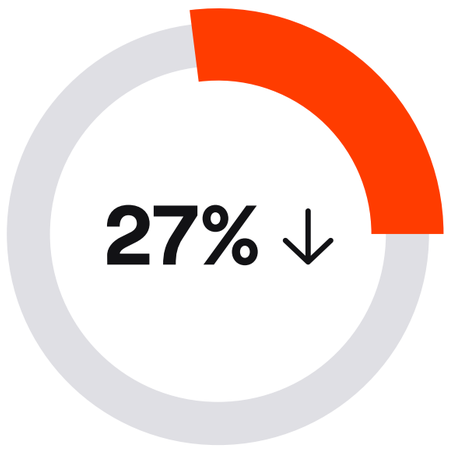

Even more telling was a **27% decrease in Lost Day Rate (LDR) (a measure of severe injuries) within that year**, indicating a substantial drop in serious incidents.

How did they achieve this? The Intenseye AI platform was configured to Amcor's specific risk profile: it monitored things like **work at height, forklift-pedestrian interactions, machine zone intrusions, and PPE compliance** in real time. Whenever an unsafe act occurred – say, a worker climbing a ladder without fall protection or someone stepping too close to a moving pallet jack – an alert would be sent to line supervisors and the individual, enabling an immediate coaching moment. Amcor's team didn't treat the system as "gotcha" surveillance; instead, they built it into a proactive safety coaching program.

According to **Amcor's multi-site EHS director**, it was "the easiest and most impactful deployment of his career" – the system seamlessly integrated with their existing camera infrastructure and began providing actionable insights from day one.

The data gathered was used in weekly safety meetings to highlight improvement areas. Heatmaps of facility risk spots allowed them to redesign problematic traffic areas (one plant re-marked walkways and added physical barriers after AI data pinpointed frequent near-hit zones). As a result, hazards that had been "invisible" or accepted as part of the job were systematically driven down. Amcor's leadership notes that this approach helped foster a more vigilant culture – employees knew that risks anywhere on the floor would be surfaced and addressed, not shrugged off.

> In essence, Amcor leveraged real-time AI as a force-multiplier for their safety team, catching problems 24/7.

The investment paid off not only in fewer injuries but also in operational continuity (less downtime) and a stronger confidence from their workforce that management truly "has their back" when it comes to safety. This case demonstrates how even a large, complex organization can rapidly improve SIF metrics by embracing Safety-III tools

Global Packaging Company

Scaling AI for SIF Prevention ROI

Case Study

ABOUT

A large paper packaging manufacturer (name withheld for confidentiality) that embarked on a global rollout of AI safety technology.

INDUSTRY

Paper Packaging

Another example comes from a company started with a pilot in 4 facilities in late 2023 and quickly saw encouraging results: sites with the AI system had a **Total Recordable Injury Rate of 1.56, compared to 2.81 at similar sites without the system – a 44% improvement** (lower is better) in recordable injuries.

Sites with Intenseye		Sites without Intenseye
1.56 Total Recordable Injury Rate	**44%** improvement	**2.81** Total Recordable Injury Rate

Following this success, they accelerated deployment to 15 plants, and then 22 plants within 9 months, noting that hazard detection accuracy remained high (83%+) and frontline teams were taking action on the insights. Across the rollout, the company observed a **48% decrease in the volume of unsafe condition alerts** – an indication that employees were adapting their behavior and conditions were improving, thanks to continuous feedback.

Dozens of proactive fixes were implemented: in the first months, over 35 safety improvements (engineering or administrative controls) were triggered by AI-identified issues, from adding guardrails in a problem area to re-training specific forklift drivers on speed control. Financially, the initiative proved its worth.

By the time they scaled to 22 facilities, **the estimated ROI was over $1.6 million** in risk reduction value (prevented incidents, reduced downtime, etc.), far exceeding the program cost.

Indeed, a cultural shift took place: sites began competing for the highest "safety scores" from the platform, and safety conversations shifted from lagging metrics to discussing real-time hazard reduction. Perhaps most importantly, this company built internal capacity around the new data – they stood up a central analytics team to review the incoming safety data weekly, identify trends across plants, and share best practices (for example, if one plant solved a frequent near-miss scenario, others learned from it).

One plant manager remarked that Intenseye became like a "second set of eyes" for supervisors – it was constantly catching things they might miss, but now those risks could be mitigated before causing harm.

The company's **Global Head of EHS** noted that since implementing AI, "hazard detection has become remarkably precise, leading to reduced injury rates and a stronger safety culture" – calling the technology "a transformative step towards modernizing EHS practices."

In less than a year, this global packaging leader demonstrated that AI-powered SIF prevention could scale rapidly without disrupting operations – in fact, operations became more efficient, as emergency stops and incident investigations dwindled. This case study highlights that with executive support and a clear strategy, an organization can integrate real-time safety tech across many sites and see measurable improvements in a short time. The experience also showed that front-line buy-in grows when workers see the system preventing accidents and not just punishing them – trust in the AI grew as it helped send people home safe day after day. For safety and operations VPs, it's a powerful example of turning cutting-edge tech into tangible safety ROI on a global scale.

Analytics & Visualizing Safety Data

From Insights to Action

One advantage of real-time SIF prevention is the wealth of data it generates. Every alert, every observation can be logged and analyzed – providing unprecedented visibility into safety performance. The challenge (and opportunity) for EHS leaders is to convert this raw data into actionable intelligence and compelling visuals that drive decision-making. In a Safety-III program, key metrics and analytics become new tools in the safety management toolbox. Here we outline a few critical analytics and how they can be visualized for maximum impact:

1 Leading Indicator Dashboards

Instead of focusing only on injury counts, leaders monitor dashboards of leading indicators – things like the number of high-risk alerts per day, average response time to safety warnings, percentage of compliance (PPE, procedures) observed, and "near-miss" counts. These can be shown as trend graphs that update in real time. For example, a chart of **daily unsafe acts observed vs. daily unsafe acts resolved** can highlight whether interventions are keeping pace with hazard occurrences. If the trend of unsafe acts is going down week-over-week, that's a positive leading indicator of risk reduction. Executives often use these dashboards as an index of safety health at a glance.

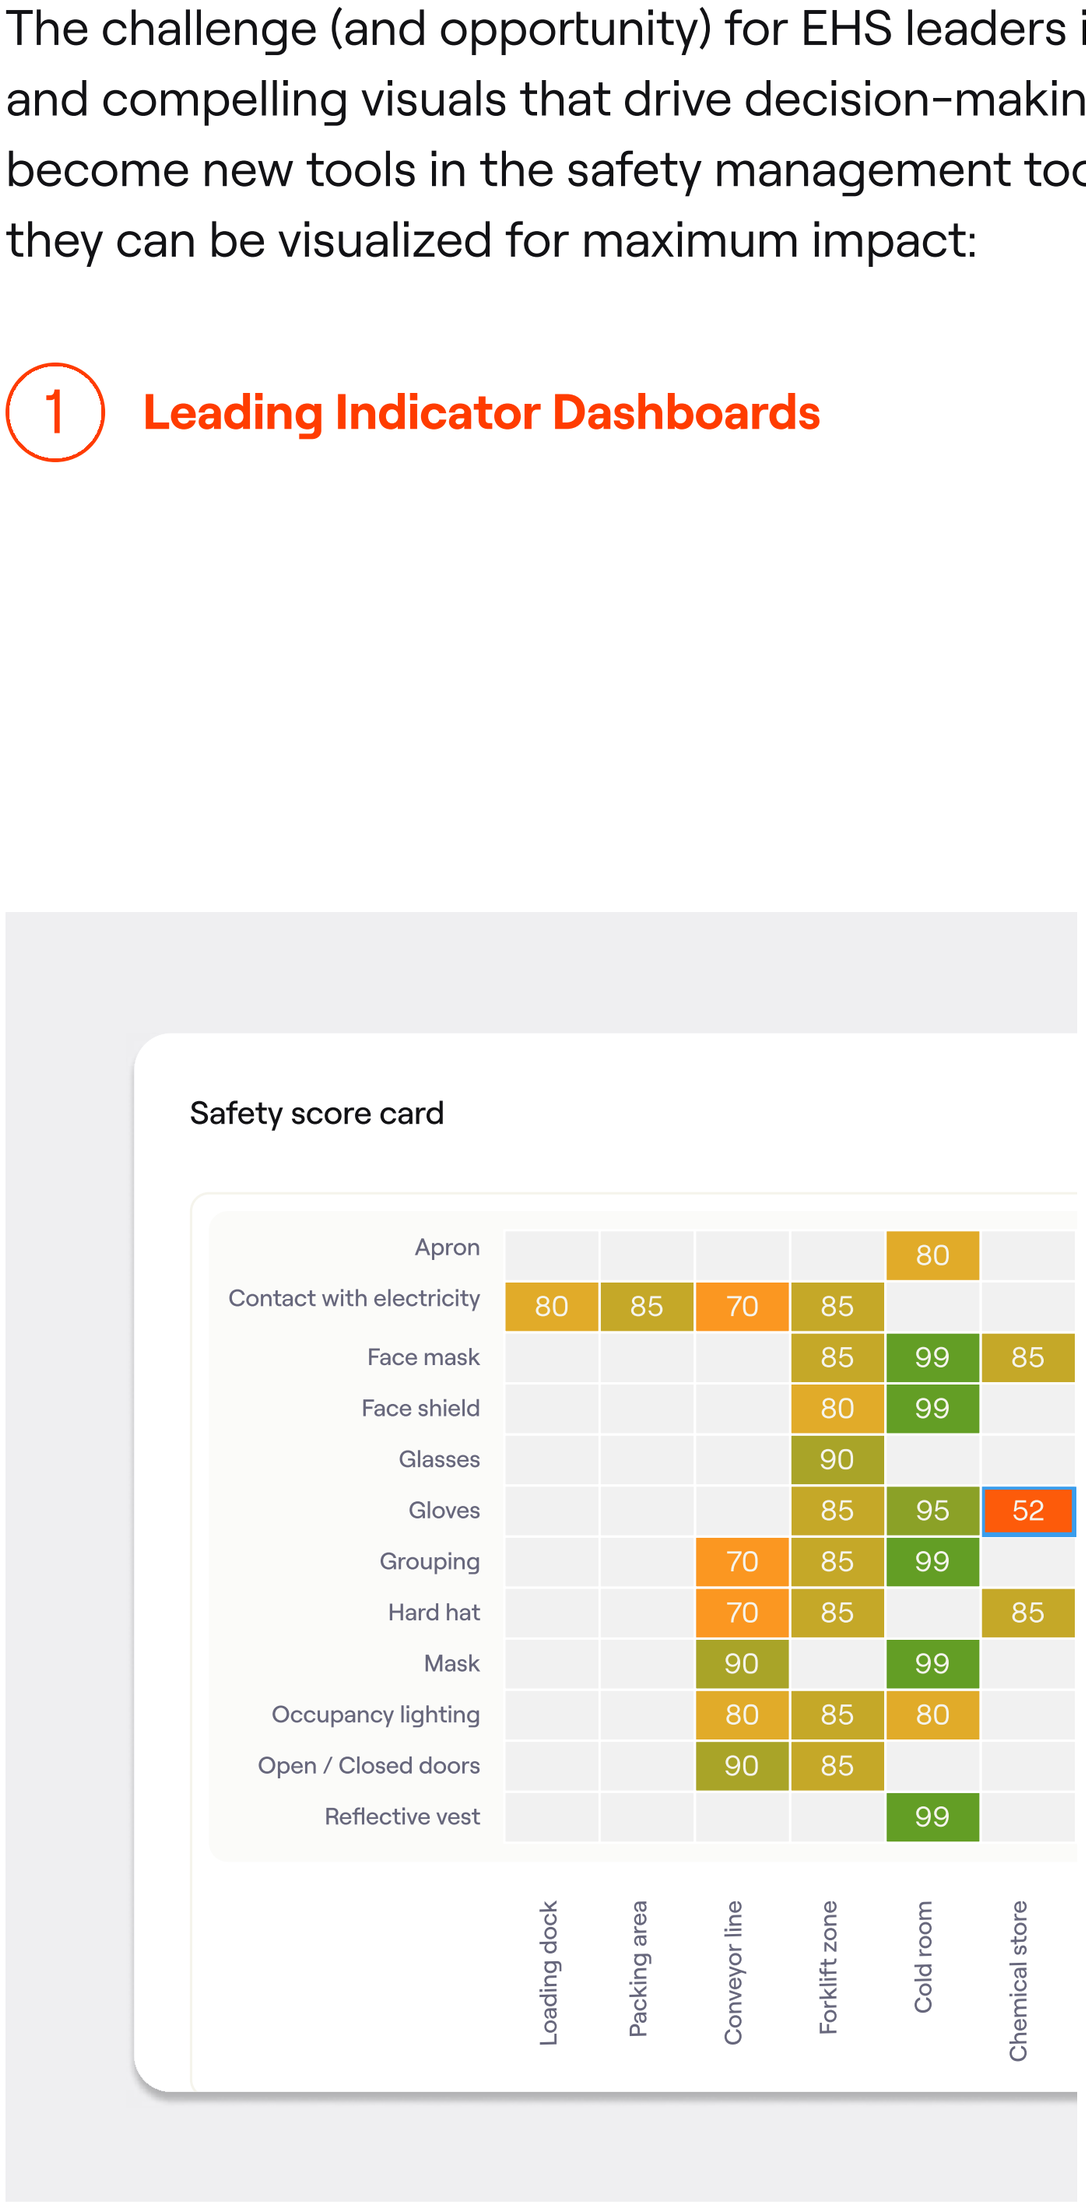

One global firm even created a composite **"Safety Score"** for each site – an index calculated from various inputs (PPE compliance %, incident severity of alerts, etc.) – and tracked it like a KPI. Sites competed to improve their Safety Scores, which drove a healthy competition and sharing of best practices. Visualizing this as a simple gauge or scorecard per site can quickly show who's leading and who might need help.

2 Heatmaps of Hazard Hotspots

Location-based analytics are extremely powerful for plant management. Intenseye and similar platforms provide dynamic, **color-coded floorplan heatmaps** that instantly reveal where the most frequent risks are appearing.

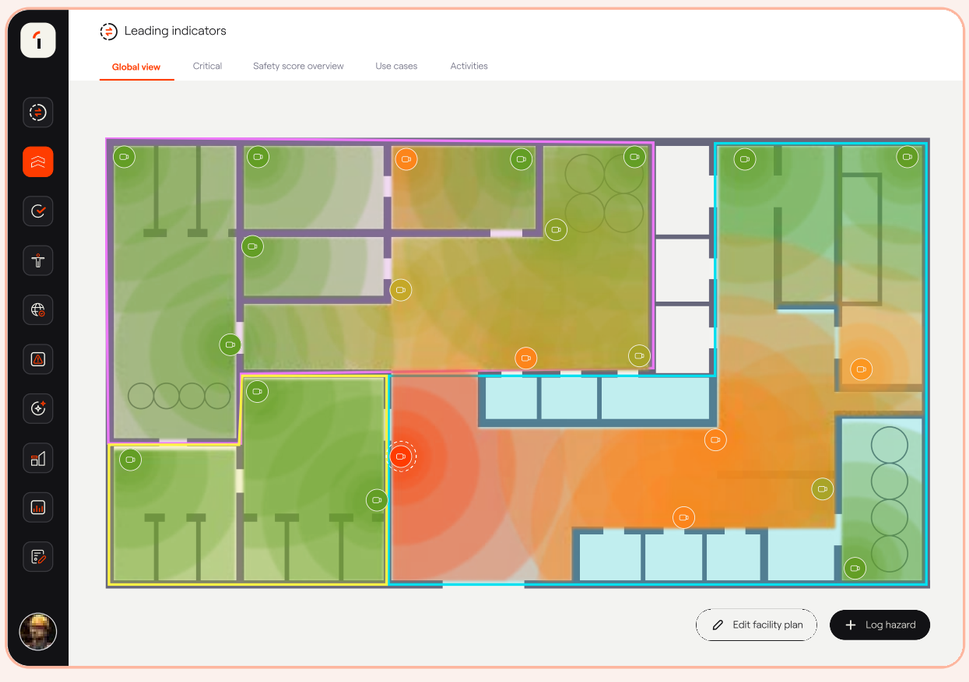

For instance, a heatmap might show concentrations of "PPE missing" alerts near a particular production line – maybe because that area has visitors who forget hardhats. Or a forklift near-miss heatmap (as discussed earlier) might highlight a dangerous intersection in a warehouse. By overlaying safety data on a map of the facility, managers can literally **see the risk landscape.** This answers the critical "where?" question in SIF prevention: where are we most likely to have the next serious incident if we don't intervene?

Such visuals guide targeted interventions – e.g. adding guardrails, mirrors, improved signage, or focused training in the hotspot areas. They also make excellent communication tools in leadership meetings (a heatmap graphic is intuitive and hard to ignore – bright red spots on your floor plan get attention).

Many companies update these heatmaps monthly to track how their risk profile changes as they implement fixes.

3 Risk Exposure Time and Cumulative Risk Metrics

Beyond counting incidents, advanced analytics calculate the exposure of workers to high risks.

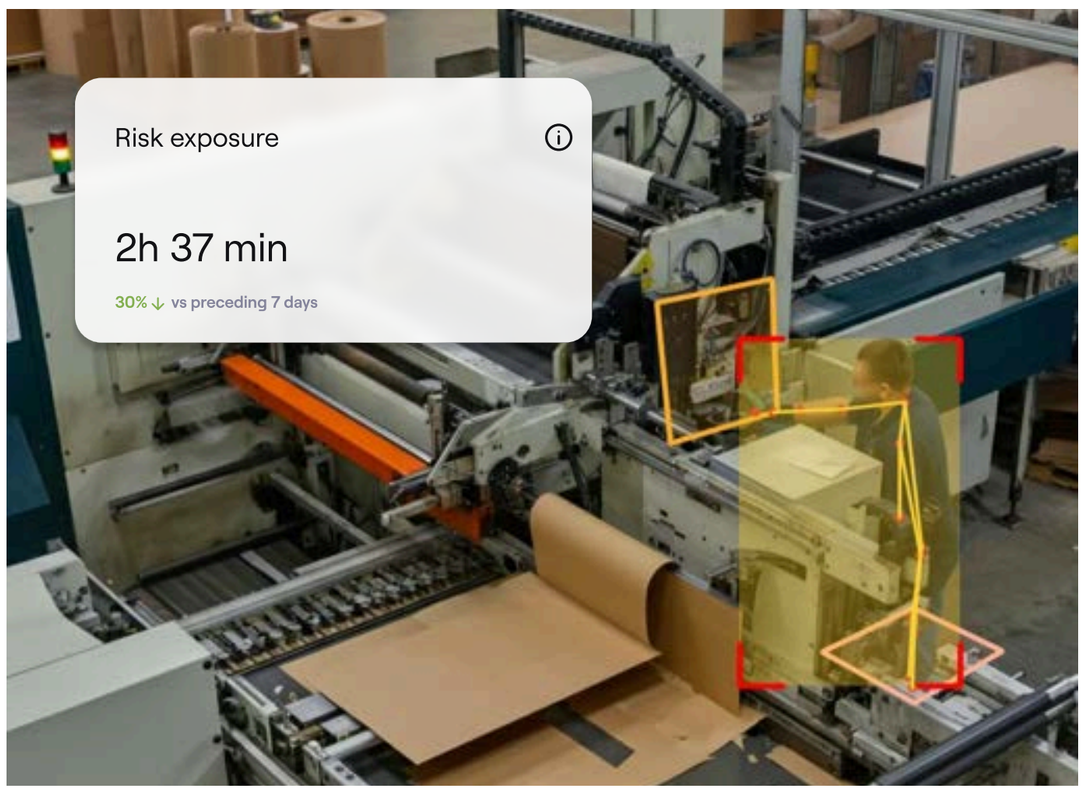

For example, Intenseye's AI-powered system can estimate how many **worker-hours** are spent in proximity to moving vehicles (a proxy for struck-by exposure) or how long machines are operated with guards open. Reducing those exposure hours is key to preventing SIFs. A visual way to present this is through **time-exposure charts** – e.g. a bar chart showing total minutes per week that any person was in a defined danger zone.

One Intenseye client tracked **"time with unsecured line-of-fire exposures"** and saw it plummet after implementing automated alerts, which was a more meaningful measure than just counting near-misses. Executives might see, for instance, that last quarter employees were exposed to 50 hours of forklift interaction time, and this quarter it's down to 30 hours after improvements – a clear reduction in SIF potential. Tables or bar graphs of these metrics can be included in quarterly reports to quantify progress in exposure reduction.

4 Alert Severity and Response Analytics

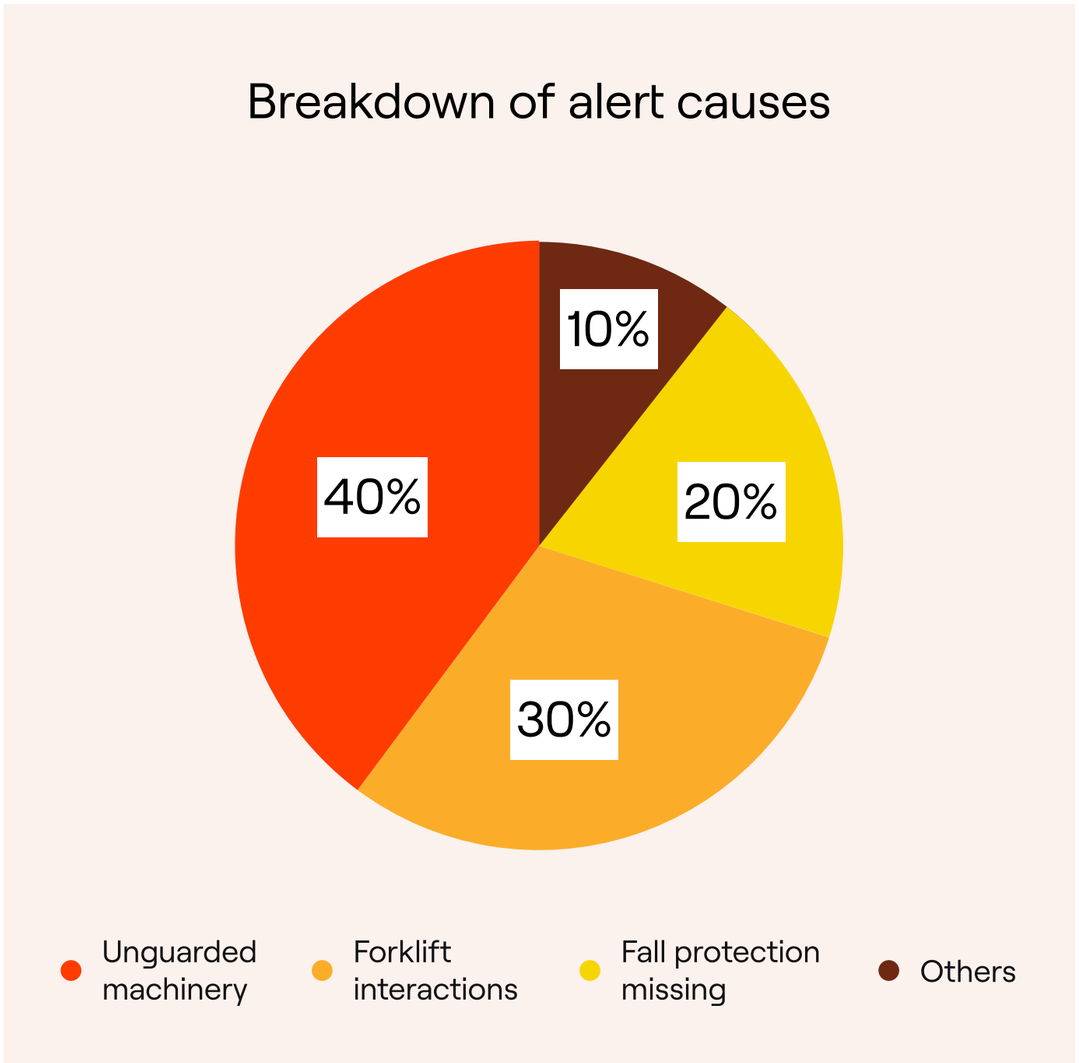

Real-time platforms log alerts by severity (high/med/low) and can track how each alert was handled. **An insightful metric is closure rate** – what percent of high-severity alerts get responded to or resolved within a target time (e.g. 1 hour). Visualizing this as a gauge or as a trend line pushes the organization to treat safety alerts with the same urgency as, say, production line downtime alerts.

You might also use pie charts to show the breakdown of alert causes: e.g. 40% of high-severity alerts last month were from "unguarded machinery," 30% from "forklift interactions," 20% from "fall protection missing," etc. This helps prioritize which risk areas to tackle first. If 40% of your serious alerts are coming from one type of hazard, that's where you focus engineering fixes or training in the next month.

5 Benchmark and Predictive Analytics

As data accumulates, companies can benchmark sites against each other and even apply predictive models. A dashboard might rank sites by SIF risk rate (perhaps normalizing for hours worked) and highlight outliers. Seeing Site A has twice the high-risk alerts per 100 workers as Site B prompts an inquiry – what is Site B doing better, or what specific challenges does Site A have? Over time, analytics might even predict which situations are most likely to lead to an incident if not corrected. For example, combining factors (like repeated medium-level PPE violations plus high workload plus new employees) might correlate with higher incident rates. AI can flag sites or departments that match those patterns so management can intervene proactively (e.g. do a safety stand-down or send additional resources before an injury happens). Presenting this in a risk matrix or bubble chart (likelihood vs severity, with bubble size = number of observations) can help convey where the next injury is most likely to occur if nothing changes.

The key with all these analytics is **turning insight into action**. Leading organizations establish routines around the data: daily production meetings that also review the prior day's safety alerts, weekly leadership calls where each plant's safety score and top issues are discussed, and monthly management reviews of trends and project statuses for risk reduction.

Some have even integrated their real-time safety data into operations control rooms – right alongside quality and production metrics – emphasizing that safety is managed with the same rigor and immediacy. By tracking the right leading indicators and visualizing them in intuitive ways, EHS and operations leaders can keep their finger on the pulse of SIF prevention. Instead of reacting to OSHA logs, they are **driving continuous improvement through data**.

One **food company's CEO** put it this way: "We went from guessing where the next accident might happen to actually knowing where our risks are daily, and that's been a game changer."

The same can be true in packaging and paper. In fact, companies that leverage these analytics report that it not only prevents injuries but also improves efficiency – because many of the risky situations (like poor ergonomics or clutter) are the same things that cause slowdowns and defects. Safety data, when used smartly, becomes business improvement data.

To support visual storytelling in your organization, consider using a mix of the above: **trend charts, heatmaps, tables, and scorecards,** enriched with real examples. For instance, show a before-and-after heatmap of a high-risk area where interventions removed the red cluster, or a trend line of declining high-risk alerts after a new training program was rolled out. These visuals make the impact of SIF prevention tangible to executives and frontline teams alike, sustaining buy-in for the program. Remember, what gets measured gets improved – and when it comes to SIF prevention, we now have more measurable leading data than ever. The companies that act on it fastest will save the most lives and gain the most value.

Actionable Recommendations

Implementing AI-Based SIF Prevention

Transitioning to a real-time, AI-enhanced safety program might sound complex, but it can be broken down into clear, practical steps. The following recommendations provide a roadmap for VPs of EHS/Ops, safety managers, and plant leaders to implement Safety-III principles in the packaging and paper industry:

1 Assess Your SIF Exposure and Set Objectives

- Begin with a frank assessment of where your serious injury/fatality risks are highest. Review past incidents, near-misses, and known hazards in your operations – do most of your SIF potentials involve machinery, forklifts, falls, or something else?
- Engage cross-functional teams (maintenance, operations, safety) to identify the top 4–5 hazardous scenarios (many likely align with those in this playbook). Set clear objectives, such as “eliminate all unsafe entries into machine X’s danger zone” or “reduce forklift-pedestrian near misses by 90%.”
- Having defined goals will guide your technology choices and configuration. Also consider what leading indicators you want to improve (e.g. increase PPE compliance rate to 99%, or respond to 100% of critical alerts within 5 minutes). These will become your performance metrics.

2 Choose Enabling Technology and Integrate it

Selecting the right **real-time safety platform** is crucial. Look for solutions that can utilize your existing infrastructure (e.g. CCTV cameras or IoT sensors you already have) to minimize deployment friction. In our industry, computer vision systems (like Intenseye) have proven effective, as they can cover wide areas and multiple hazard types at once.

Work closely with your IT and OT (operations technology) departments to ensure the system can be integrated securely with your networks and doesn't interfere with control systems.

Start small if needed – for example, outfit one or two high-risk areas with the AI cameras and sensor integrations. This might be the corrugator and forklift loading zone, for instance.

Ensure the system is calibrated for your environment: correct camera angles, appropriate alert thresholds (you might set a forklift alert distance of, say, 3 meters to start, then adjust). Integration also means connecting the alert system to the right channels – whether that's SMS/email to supervisors, flashing stack lights on the factory floor, or notifications in a safety management app.

Aim to make alerts impossible to ignore but also not overwhelming (it's okay to start with only the most critical 5–6 rule violations being monitored, then expand later). In parallel, plan for how this will tie into existing systems: you may want to feed the data into your EHS software or incident database so all information is in one place. Many AI platforms offer open APIs for integration – use them to avoid creating data silos.

3 Engage and Train Your People

Introducing AI monitoring in the workplace can raise concerns (privacy, "big brother", job implications), so it's vital to get buy-in from day one. Communicate transparently with employees and supervisors about what the system does and why you're implementing it – emphasize that the goal is to keep everyone safe and to prevent tragedies, not to punish minor mistakes. In fact, highlight that the focus will be on **coaching not discipline for most alerts**. Provide training sessions or demos so workers see how it works (for example, show a live or recorded example of the system catching an unsafe act and the kind of alert it sends).

Many will be amazed that technology can even do this. Address questions: "What if it flags me when I had to do something for production?" – make sure they know there will be a human review and it's there to prompt safer alternatives, not to automatically blame. Also, train your supervisors and EHS staff on how to respond to alerts.

Define protocols: e.g. "If a high-severity alert comes in, immediately stop the task and correct it, then log the corrective action." Teach them to use the analytics dashboard – they should know how to pull up the week's trends, acknowledge alerts in the system, and input notes (many platforms let you mark an alert as 'addressed' with a comment like "employee retrained on PPE"). Essentially, integrate the system into your daily safety management. Some companies assign a dedicated "Safety AI champion" at each plant to ensure it's being utilized and to gather feedback from operators.

Early involvement of employees in system roll-out (even allowing them to suggest what hazards to monitor) can turn skeptics into advocates. As they start seeing the AI catch real dangers, trust will build. Reinforce positive behavior: when the system reports improvements (e.g. 3 weeks with no one entering zone X unsafely), celebrate that with the team. This human element of change management is as important as the tech itself.

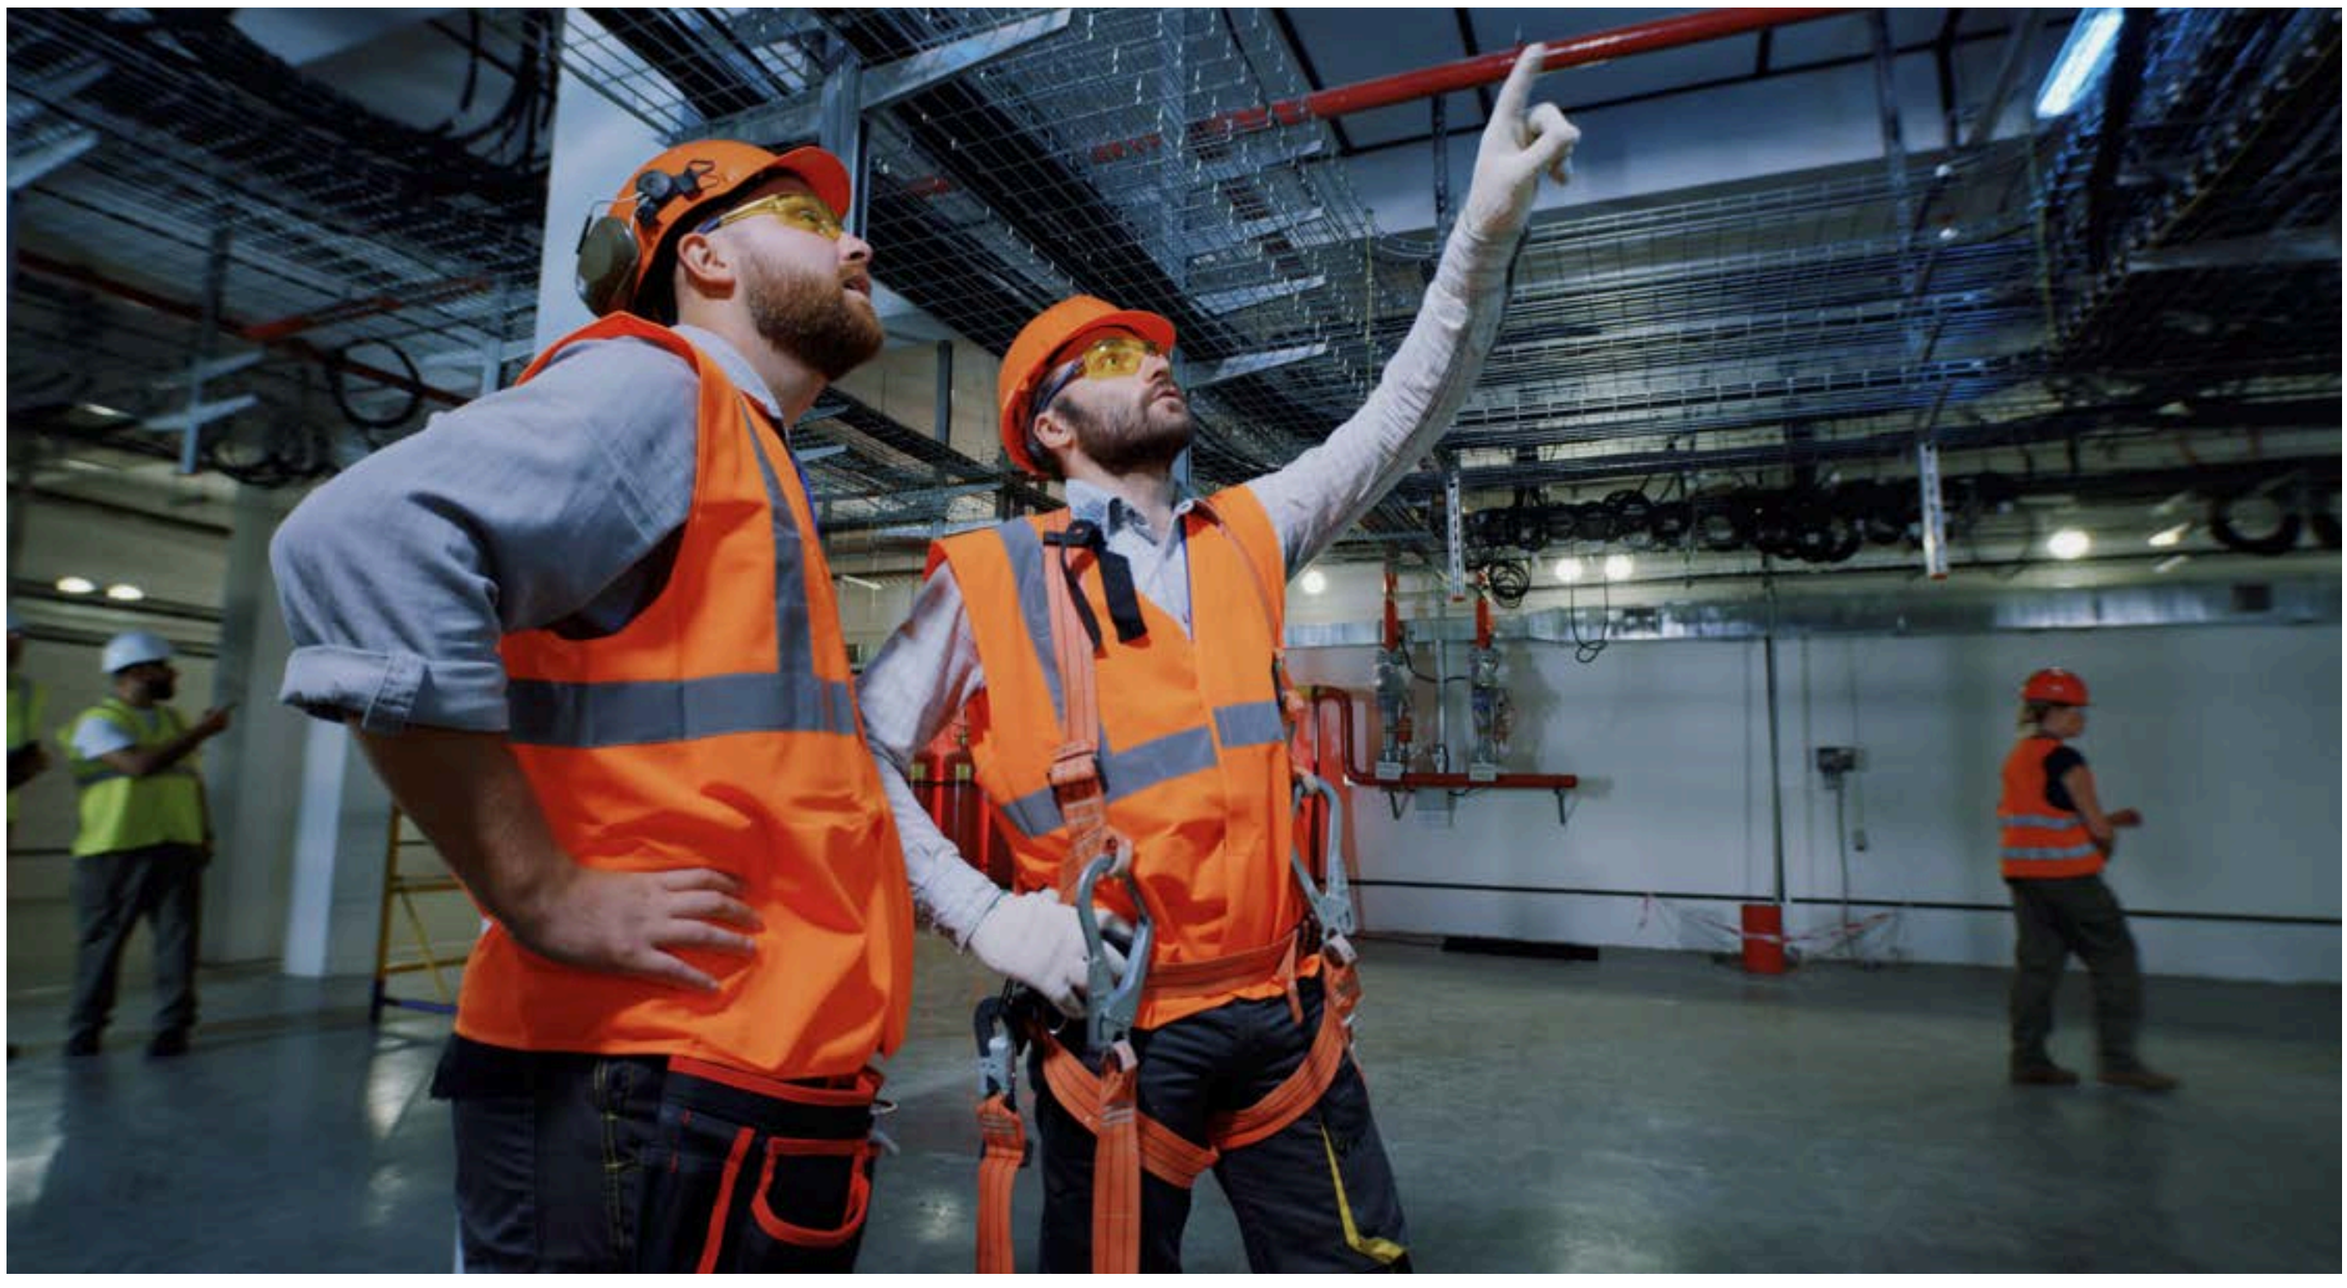

4 Realign Safety Metrics and KPIs

To fully adopt a proactive Safety-III approach, you should update the way you measure and incentivize safety performance. Incorporate leading indicators from the AI system into your regular KPIs. For example, you might add metrics like “Unsafe situations resolved per month”, “Average response time to critical alerts”, or “% of inspections completed via AI data vs target.” This doesn’t mean abandoning lagging metrics like TRIR – but it means management talks about both now.

For instance, in monthly ops reviews, present not just the injury count, but also how many high-risk events were averted. Some organizations have tied managerial bonuses or performance evals in part to these leading safety metrics (e.g. bonus criteria might include maintaining a Safety Score above a certain level, or closing 95% of safety alerts within 24 hours). Such alignment sends a clear message: preventing the potential incident is as valued as preventing the actual incident. It drives home that near misses and hazards are treated with seriousness.

When VPs ask their plant managers, “We had zero injuries this quarter, but what is our exposure profile? What are the top things the AI is catching and are we fixing them?” – it signals a cultural shift. We also recommend establishing a regular cadence for reviewing analytics (as noted in the prior section). Perhaps you form a Safety Analytics Review Board that meets quarterly to look at company-wide data and make recommendations (much like quality or production continuous improvement committees).

Over time, track how your proactive metrics correlate with lagging outcomes – you will likely find that as your leading indicators improve, your injury rates follow. This data-driven validation will help secure ongoing investment in the program. In short, update your scorecards to **include Safety-III metrics** so that everyone from C-suite to shift leads is accountable not just for “what went wrong” but for “what we prevented from going wrong.”

5 Pilot, Learn, and Scale Up Strategically

Start with a manageable scope, prove the concept, then expand. A common approach is a **pilot in a single facility (or a focused area)** for 3–6 months.

- Pick a site with engaged leadership and a genuine need (e.g. a plant that has seen a couple of recent serious near-misses – they'll be highly motivated).

- Use the pilot to refine the system configuration (you might discover, for instance, that the AI is overly sensitive in one area and adjust the threshold to reduce false alarms).

- Collect baseline data and improvements – like how many unsafe acts were identified and corrected, or any reduction in incident rates compared to the prior period. Also gather feedback from users: did the alerts make sense? Were any nuisance alarms?

- Use these insights to tweak your approach. Then build the business case for scaling: calculate the pilot's ROI in terms of incidents prevented or productivity gained (e.g. one pilot site saw a 44% lower TRI than sister sites without the system).

- Present these wins to senior management to secure buy-in (if you're the VP, use it to convince the Board or Finance). When scaling, take a phased approach: for example, move from 4 pilot sites to 10 more sites in the next phase (an "accelerate" phase), before full rollout to all 20+ sites ("expand"). This phased scale allows adaptation to each site's layout and culture.

- Leverage internal champions – perhaps the safety manager from the pilot site can mentor the new sites.

- Ensure each site getting the technology has a proper kickoff and training similar to the pilot. It can be valuable to share data across sites: create internal case studies or newsletters ("Plant A reduced forklift speeding by 80% using the new system – here's how they did it"). This fosters a positive peer pressure and knowledge transfer. Scaling globally in under a year is feasible – one packaging company went from 4 to 22 facilities in 9 months – but only with strong project management and executive support to remove roadblocks.

- Monitor the implementation closely; if one location is lagging (e.g. not using the system fully), visit them, understand issues, and re-engage. Remember, technology is only as effective as the process around it. So standardize new procedures (like how to handle an AI-detected violation) across sites as you expand.

> By the end of the scale-up, you should embed the AI platform into the fabric of your safety management system company-wide.

6 Sustain and Continuously Improve

Implementing the system is not a one-and-done project – it's the start of a continuous improvement cycle. Treat the AI like a member of your safety team that needs periodic evaluation and tuning. Set aside time (perhaps bi-annually) to review the rules/algorithms you're using: do you need to add new ones as processes change? Are there false positives that can be reduced by refining the AI model or adjusting camera positions?

Engage with your technology provider for updates – AI models improve, and new features (maybe new hazard detection capabilities) are often added. Take advantage of those to expand your SIF prevention arsenal. Also, rotate focus when needed: if one hazard category is well controlled now (say, you've achieved 99% hard-hat compliance), you might shift emphasis to another emerging issue (maybe ergonomics or quality-related safety issues).

Use your analytics to identify these opportunities. Culturally, continue reinforcing the value of proactive safety.

Share success stories: for example, if a machine guard alert literally saved someone's life by stopping a machine, make sure every employee hears about that in safety meetings – it reinforces why the system exists. Likewise, be transparent about trends: "Last month we had 5 high-severity alerts – let's get that to zero this month by focusing on X." Many leading companies form cross-plant safety committees that specifically discuss SIF precursors (often leveraging the AI data as their agenda). This keeps the momentum going and avoids complacency.

Over a year or two, these practices can truly transform your safety trajectory – one company saw their injury rates drop consistently each quarter as they iteratively fixed issues flagged by the system.

Finally, don't forget to celebrate the wins: fewer incidents, compliance milestones, people whose quick action on an alert prevented harm – recognize and reward these. It shows that real-time prevention is now part of "the way we work" and that everyone's contributions to it matter.

By following these steps – assessing risks, deploying the right tech, training people, aligning metrics, scaling up thoughtfully, and continuously improving – an organization can successfully integrate AI-based, real-time SIF prevention into its operations. The process does require investment and change management, but the payoff is huge: literally saving lives, preventing debilitating injuries, and fostering a culture where safety is truly ingrained in every action.

As you implement, keep the critical question in mind: "How will this help us spot and stop the next potential fatality?" If each step you take can clearly answer that, you're on the right track.

Conclusion

A New Era of Proactive Safety Leadership

The containers, packaging, and paper industry stands at a pivotal moment in safety. The data is clear – we have plateaued in traditional safety improvements, and in some cases serious incidents are creeping upward. But as this playbook has illustrated, we also have unprecedented tools at our disposal to change that trajectory.

Real-time, AI-driven SIF prevention is not science fiction; it's here and now, already delivering measurable results for companies that have embraced it. Adopting these technologies and approaches is more than a technical upgrade – it is a moral and business imperative. Every executive in this industry knows the feeling of dread when the phone rings with news of a serious accident. By investing in Safety-III systems, we are saying no to accepting "freak accidents" as fate. We're equipping our teams with an always-on safety net, and in doing so, we're fulfilling the most fundamental duty of leadership: protecting our people.

The benefits extend beyond the immediate prevention of harm (as if that weren't justification enough). Companies that lead on safety tend to lead in operational excellence as well. Fewer incidents mean less downtime, less disruption, and more consistent productivity. Near misses addressed in real time often reveal process weaknesses that, when fixed, also improve efficiency. A safer operation is often a better-run operation – housekeeping, maintenance, and training all tend to be sharper in a workplace that prioritizes serious injury prevention. Additionally, top-tier safety performance has become a **competitive differentiator.** Customers and partners are increasingly looking at ESG and safety records. A reputation for cutting-edge safety can enhance brand value and stakeholder trust. And consider workforce aspects: in an era of labor shortages, employees want to work where they feel safe. Demonstrating that you have an AI "guardian" watching out for them, and that you truly care about sending them home unharmed, helps attract and retain talent. It's part of being an employer of choice.

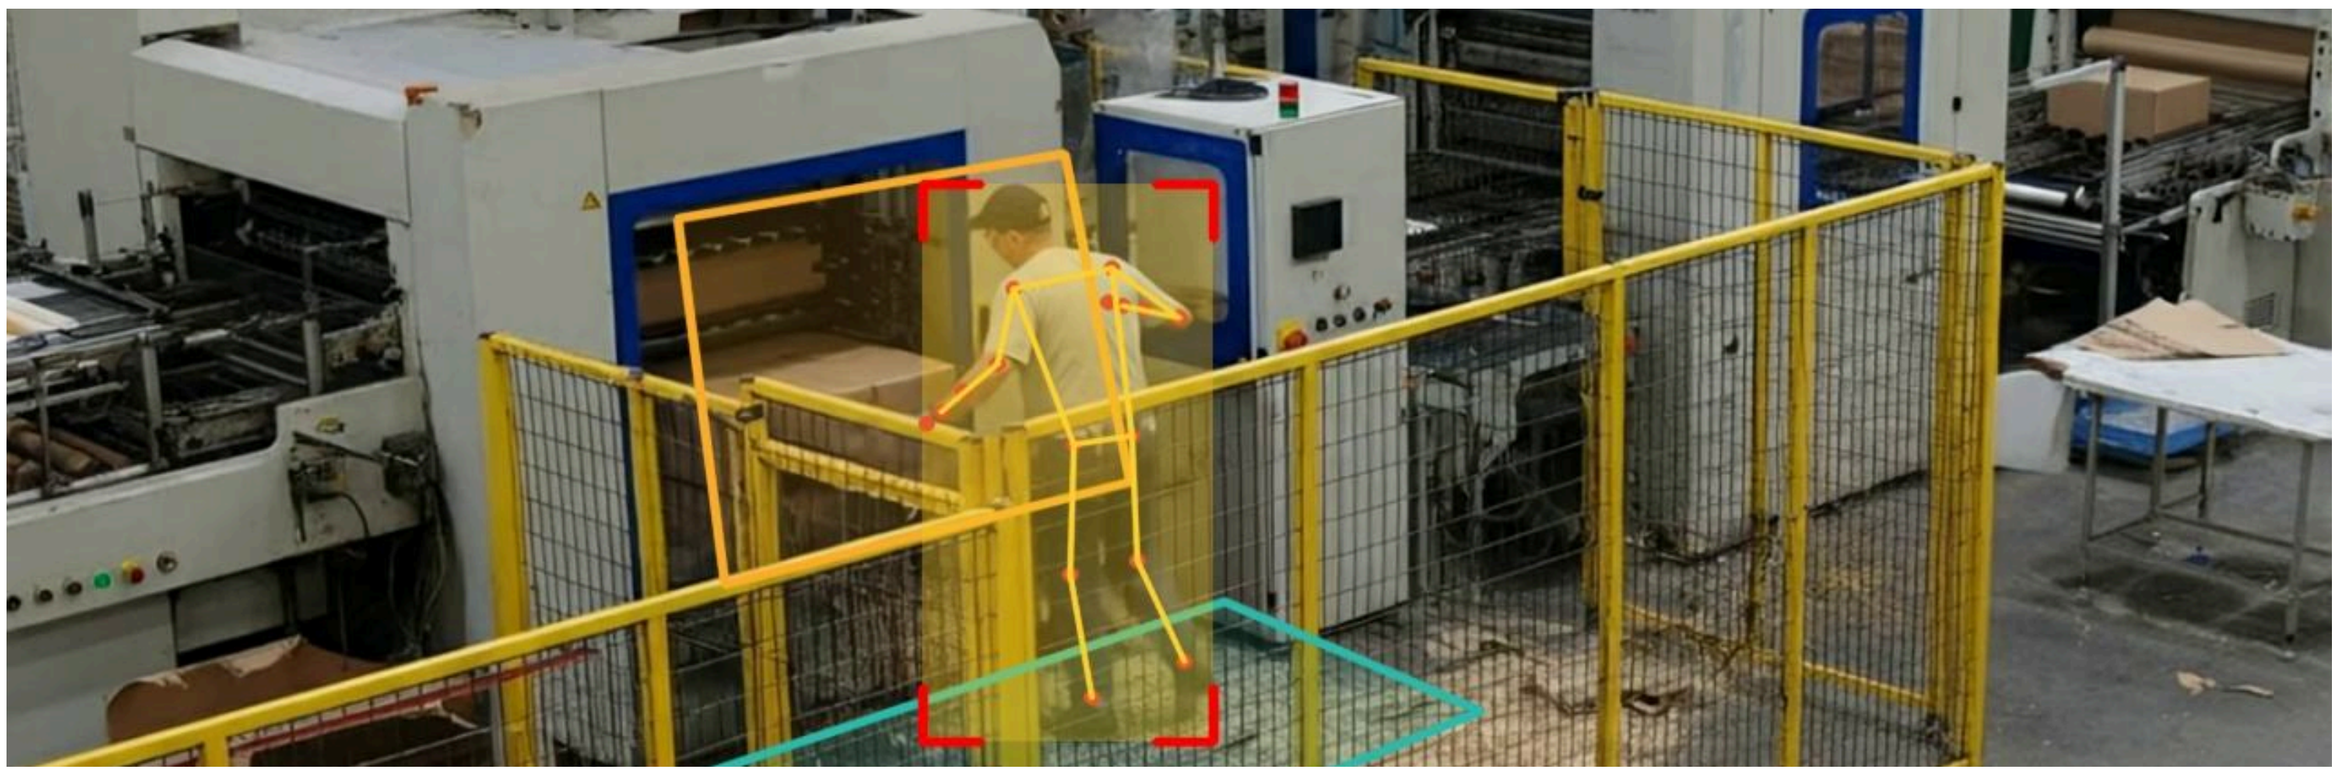

We must also acknowledge the broader industry commitment. Organizations like AF&PA have declared goals of **zero workplace injuries by 2030.** That is a bold target that cannot be met with complacency or doing "more of the same." It will require innovation and a willingness to leap forward. The trailblazers referenced in this playbook – from Amcor to the anonymized case and others – have shown that **serious injury and fatality risks can be dramatically reduced,** and importantly, **this can be achieved without hindering production or adding burdensome work**. In fact, in those examples, operations became smoother and more efficient as emergencies and accidents dwindled. They have effectively paved the path and worked out many initial kinks. Now it's up to the rest of the industry's leaders to follow suit and scale these successes.

In closing, the call to action for packaging and paper executives and safety professionals is clear: **take the leap into real-time SIF prevention**. Equip your plants and people with the technology and processes that can detect danger the instant it arises. Foster a culture that values foresight as much as hindsight. Use data and AI as allies in the fight against workplace tragedy. By doing so, you will not only prevent the next potential fatality – you will exemplify the kind of forward-thinking leadership that drives sustainable success. The tools are ready, the business case is compelling, and – most importantly – **the human case is unquestionable**. Every worker deserves to return home safely. With real-time SIF prevention, we can honor that promise better than ever before. The companies that act now will not only save lives but also lead the industry into a new era where serious injuries and fatalities are no longer seen as inevitable, but as preventable. It's both the right thing to do and the smart thing to do. The future of safety in packaging and paper is being written today – let's write it together, with vision and determination, one prevented incident at a time.