

AI, Due Diligence, and the False Comfort of Fast Compliance



There's a reason AI is spreading quickly through compliance work.

It's fast.

A safety manager can generate a draft policy in seconds. A supervisor can turn a procedure into a toolbox talk before the morning meeting. HR can ask for a harassment training outline. Operations can request a checklist for contractors. A training coordinator can create quiz questions, refresher content, sign-in sheets, learning objectives, and email reminders without starting from a blank page.

For overloaded safety teams, that feels like relief. In many organizations, safety leaders are under constant pressure to produce more with less. More orientations. More refresher training. More documented procedures. More corrective actions. More compliance updates. More proof for audits, insurers, customers, prime contractors, regulators, and senior leadership.

AI appears to solve one of the biggest operational frustrations in safety: the amount of time it takes to produce the material.

But that's also where the risk begins.

Fast compliance can feel like real compliance before it has earned that status. A document can be generated quickly. A training module can be assigned quickly. A checklist can be distributed quickly. A record can be created quickly. None of that proves the organization understood the legal duty, identified the actual hazard, selected the right control, trained the right people, verified competence, or retained the kind of evidence that matters after something goes wrong.

That's the false comfort of fast compliance.

The FDA's 2026 warning letter to Purolea Cosmetics Lab offers a useful warning for every regulated industry. FDA said the company used AI to create specifications, procedures, and production or control records, but failed to review those AI-generated documents to ensure they were accurate and compliant. FDA also identified overreliance on AI as part of the inspection findings. The key lesson is simple: regulators may tolerate AI-assisted work, but they won't excuse uncontrolled AI-assisted compliance failure.

Workplace safety leaders should assume the same principle applies to them.

If an employer uses AI to produce a lockout procedure, fall protection talk, confined space checklist, forklift refresher, WHMIS/HazCom handout, incident investigation

guide, or supervisor safety script, the employer still owns the result. AI doesn't become the employer. AI doesn't become the competent person. AI doesn't become the safety committee. AI doesn't become the supervisor. AI doesn't become the legal reviewer. AI doesn't become the trainer of record. It's a tool, not a duty holder.

That distinction matters because due diligence is not about whether a company had paperwork. It's about whether the company took reasonable precautions under the circumstances. CCOHS describes due diligence in OHS as the reasonable steps taken to prevent workplace harm, including identifying hazards and taking corrective action. OSHA's training materials also make clear that training obligations are embedded across many standards, and that guidance does not replace the employer's actual compliance duties under applicable standards.

That's why AI misuse in safety is more than a bad content problem. It can become evidence of a weak safety management system.

Consider a common scenario.

A construction company is expanding quickly. It has projects in multiple jurisdictions, a mix of employees and subcontractors, and a safety department that's struggling to keep up. The company decides to use AI to modernize its safety documentation. It generates new toolbox talks, inspection forms, pre-task planning templates, incident investigation questions, orientation materials, and supervisor guides. Everything looks cleaner than the old documents. The documents are easier to read. Supervisors like them because they're short. Senior management likes them because the rollout looks efficient.

For several months, the system feels better. Then there's a serious fall.

The investigation turns up uncomfortable questions. The fall protection talk used the wrong trigger height for the jurisdiction. The rescue section was generic. The site-specific fall protection plan had not been updated after work conditions changed. Supervisors believed workers had been trained because the AI-assisted refresher was assigned through the LMS, but nobody verified that the workers understood the anchor points, travel restraint limitations, inspection steps, or rescue expectations for that site. The company had records, but the records showed that workers completed training that didn't match the actual work.

The problem wasn't that AI had been used.

The problem was that AI had entered the safety system without enough control.

From an insurance and loss prevention perspective, that's a serious concern. Insurers, claims teams, and defence counsel aren't only interested in whether a document exists. They're interested in whether the document helps prove the employer had a functioning system. **After a serious injury, the useful questions are more demanding.**

1. Was the procedure accurate?
2. Was it current?
3. Was it site-specific?
4. Was it reviewed by someone competent?
5. Was it consistent with the applicable law?
6. Was it consistent with manufacturer instructions?
7. Was it communicated to affected workers before exposure?
8. Was the worker trained in a language and format they could understand?
9. Was understanding tested?
10. Was performance observed?
11. Were supervisors trained to enforce it?
12. Were deficiencies corrected?
13. Was the document updated after incidents, inspections, equipment changes, or

regulatory changes?

That's the difference between having material and having proof.

AI tends to make material easier. It doesn't automatically make proof stronger.

In fact, AI can weaken proof if the organization can't explain how the material was created, reviewed, approved, assigned, and validated. Imagine sitting across from an inspector, claims investigator, plaintiff's lawyer, union representative, or corporate customer after a serious incident. The employer presents a safety procedure. The next question is obvious: "Who wrote this, and how did you confirm it was right?"

If the real answer is "someone generated it with AI and we cleaned it up," the organization has a problem.

Not because AI was involved. **Because the review process was inadequate.**

The risk is even higher in North America because OHS compliance is fragmented. A single organization may operate under federal OSHA, state-plan OSHA, Canadian provincial or territorial OHS laws, federally regulated Canadian requirements, sector-specific rules, workers' compensation board expectations, customer prequalification systems, insurance loss control recommendations, consensus standards, and internal corporate policies. Training obligations and procedural expectations can vary by jurisdiction, industry, hazard, and worker role.

AI can flatten those differences. It can produce one smooth answer where the law requires several precise ones.

That's especially dangerous for topics such as fall protection, lockout/tagout, confined spaces, first aid, heat stress, workplace violence, harassment, WHMIS/HazCom, respiratory protection, powered industrial trucks, machine guarding, construction orientations, young worker training, supervisor competency, and JHSC or safety committee obligations. These areas often involve specific definitions, thresholds, procedures, records, refresher expectations, competent person duties, or jurisdictional variations.

A generic answer may be useful as a starting point. It should not be treated as the final control.

The other problem is that AI can make administrative controls look stronger than they are. Training, procedures, signs, reminders, checklists, and policies are important, but they often sit lower in the hierarchy of controls than elimination, substitution, and engineering controls. NIOSH describes the hierarchy as a framework for reducing or removing hazards, with elimination, substitution, engineering controls, administrative controls, and PPE ranked in that order of general effectiveness. CCOHS describes the hierarchy similarly as a step-by-step approach for eliminating or reducing workplace hazards, ranked from most effective to least effective.

That matters because AI is very good at producing administrative controls.

It can write a policy. It can draft a reminder. It can create a checklist. It can generate a training script. It can produce a toolbox talk. It can suggest PPE. But if the hazard should be eliminated, engineered out, isolated, guarded, redesigned, or controlled through a physical change, AI-generated training may create the illusion of action while leaving the hazard in place.

That's where loss prevention professionals should pay close attention.

If a warehouse has repeated struck-by near misses because pedestrians and forklifts share blind intersections, a better safety talk might help. But the stronger control may be traffic redesign, barriers, mirrors, warning systems, scheduling changes,

designated walkways, speed controls, or separation of people and equipment.

If workers are repeatedly exposed to heat stress because production targets make breaks unrealistic, another heat stress refresher may not address the real problem. The employer may need workload changes, acclimatization planning, shaded rest areas, cooling measures, supervisory triggers, medical response planning, or changes to work scheduling.

If workers bypass guards because the machine jams repeatedly, another reminder not to bypass guards may be a weak response. The better question is why the jam occurs, whether the machine can be modified, whether guarding can be improved, whether maintenance is adequate, and whether lockout is practical and enforced.

AI-generated content can support those discussions. It can't substitute for them.

The biggest executive risk is that AI may allow senior leaders to believe the organization is more compliant than it is. A dashboard can show high completion rates. A folder can contain updated policies. A training library can appear extensive. A safety portal can look modern. But if the underlying content hasn't been validated, if the training doesn't match the work, if supervisors aren't reinforcing it, and if hazards remain uncontrolled, the organization may have improved the appearance of compliance without improving the substance.

That's dangerous because appearance often satisfies management until an incident tests the system.

After a serious event, fast compliance doesn't matter much. Defensible compliance does.

Defensible compliance means the employer can show a logical chain of action. The hazard was identified. The legal and operational requirements were understood. Controls were selected according to risk. Workers were trained before exposure. Supervisors were prepared to enforce the rules. Competence was verified where needed. Records were retained. Changes were made when conditions changed. The system was audited. Deficiencies were corrected.

AI can help organize that chain. It can't create it on its own.

So what should organizations do?

The answer isn't to ban AI from safety. That would be unrealistic. It would also miss the opportunity. Used properly, AI can help safety teams simplify dense procedures, draft first versions, convert technical content into supervisor talks, create scenario-based exercises, identify gaps in a training outline, improve readability, support multilingual communication, and reduce administrative load.

But AI needs governance.

The first step is to define where AI can and can't be used. Low-risk use cases might include drafting plain-language reminders, creating first drafts of toolbox talks, summarizing internal procedures for review, brainstorming scenario questions, or converting approved content into shorter learning formats. Higher-risk use cases include legal interpretation, hazard assessments, critical procedures, incident findings, root cause analysis, disciplinary recommendations, medical or exposure advice, emergency plans, and anything connected to high-risk work.

The second step is to require human ownership. Every AI-assisted safety document should have an accountable owner. That owner should know what the document is for, what sources were checked, who reviewed it, when it was approved, and when it needs to be updated.

The third step is to require competent review for high-risk content. A competent

reviewer may be an OHS professional, supervisor, engineer, maintenance lead, HR specialist, legal reviewer, JHSC member, external consultant, or another qualified person depending on the issue. The point is not job title. The point is competence.

The fourth step is to require source validation. AI-generated compliance content should be checked against the actual law, regulator guidance, recognized standards, manufacturer instructions, internal procedures, incident history, and site-specific conditions. If the output can't be traced back to reliable sources, it shouldn't be treated as compliance-ready.

The fifth step is to separate drafting from approval. AI can help draft. It should not approve. The approval process should be visible, documented, and proportionate to the risk.

The sixth step is to control versions. Organizations should know which version was used, when it was assigned, who completed it, what changed, and who approved the change. This is especially important when training is used across multiple locations or jurisdictions.

The seventh step is to test understanding and application. For high-risk tasks, a quiz isn't enough. The employer may need supervisor observation, field demonstration, scenario discussion, practical evaluation, coaching, or retraining after deficiencies.

The eighth step is to audit AI-assisted materials periodically. Laws change. Equipment changes. Work changes. A document that was reviewed last year may not remain reliable this year. This is particularly important for multi-jurisdiction employers.

The ninth step is to ask vendors direct questions. If a training provider, consultant, compliance platform, translation service, or content vendor uses AI, the employer should understand how outputs are reviewed, updated, controlled, and documented. The vendor's use of AI doesn't remove the employer's responsibility, but a strong vendor process can reduce risk.

The tenth step is to connect AI governance to the broader safety management system. AI shouldn't be a side issue handled informally by whoever creates content. It should be part of document control, training governance, incident response, management review, procurement, and audit processes.

This is where a platform like SafetyNow can be positioned carefully and credibly. The market doesn't need more unsupported compliance content. It needs better systems for delivering reviewed training, assigning it consistently, tracking completion, maintaining records, reinforcing learning, and giving organizations practical evidence that training occurred. AI may help with speed and customization, but the real business value is still control, consistency, cognitive learning, and defensible proof.

That's the message executives and insurers need to hear. AI is not a compliance strategy. It's a production tool inside a compliance strategy.

Used casually, it creates risk. Used carefully, it can reduce administrative friction without weakening accountability.

For insurers and workers' compensation partners, this issue is especially relevant because training quality affects loss prevention, claim defensibility, and employer behaviour. If policyholders use AI to generate weak training, they may believe they've reduced risk when they've only created a better-looking file. If they use controlled training systems, reviewed content, supervisor reinforcement, and documented competency checks, they're more likely to reduce real exposure.

The business case isn't anti-AI. It's anti-false confidence.

Safety leaders should use AI where it helps, but they should be ruthless about the boundary between assistance and accountability. AI can suggest. AI can draft. AI can summarize. AI can translate. AI can reformat. AI can help create scenarios.

- But AI can't walk the floor.
- AI can't inspect the machine.
- AI can't see the blind corner.
- AI can't know that the crew skips the procedure when production runs late.
- AI can't tell whether a new worker understood the instruction.
- AI can't verify that a supervisor corrected the unsafe shortcut.
- AI can't prove due diligence after a worker is injured.

That still belongs to the employer.

The organizations that understand this will be able to use AI without letting it hollow out their safety systems. They'll move faster, but not blindly. They'll produce better materials, but still review them. They'll train more consistently, but still verify application. They'll use records, but not hide behind them.

The organizations that don't understand this will create beautiful compliance theatre.

They'll have documents, dashboards, completion rates, and polished language. But when something goes wrong, the question won't be how quickly the training was created. It'll be whether it was accurate, specific, reviewed, understood, applied, and supported by real controls.

Fast compliance feels good.

Defensible compliance protects the organization.

More importantly, it protects the worker.