AI Use Policy, Training, and Governance

Prepared for NAED Digital Center of Excellence publication and member guidance

Scott A. Wagner | Director of Industry Transformation | NAED | swagner@naed.org

Executive Summary

Building a effective AI usage policy is becoming one of the first real tests of AI maturity in electrical distribution. Many distributors have moved beyond the question of whether employees will use AI. They already are. Employees are using AI to draft emails, summarize meetings, compare supplier information, analyze spreadsheets, prepare customer responses, generate marketing content, support HR communication, and simple exploration of AI-enabled tools like ChatGPT and Microsoft Copilot. The activity is real. The operating discipline is often still catching up.

The greatest policy risk is not that a distributor will have no written policy. The greater risk is that the policy will be too vague, too long, too legalistic, or too disconnected from daily work to shape behavior. A policy that employees cannot remember in the moment will not protect customer information, employee data, pricing logic, supplier agreements, margin strategy, contracts, or operational commitments. It will create the appearance of control while leaving the real decisions to individual interpretation.

A practical AI use policy should be treated as an operating tool. It should help responsible people move with confidence. It should answer four questions clearly:

- Which tools are approved?
- What data can and cannot be used?
- Where is human review required?
- Who employees should ask when they are unsure?

The policy should not try to predict every edge case. It should create enough clarity that employees and managers know the safe path before pressure pushes them into unguided experimentation (also known as, “shadow AI”).

Safe AI training should be approached in very much the same way a company approaches email phishing or cybersecurity training. One announcement will not change behavior. One policy document will not stop unsafe usage. Employees need repeated, role-based, scenario-driven training that reflects the work they actually do in branches, sales, purchasing, warehousing, HR, finance, marketing, credit, and operations. They need to learn what not to paste into a tool, how to challenge output, how to identify work slop, and how to escalate uncertainty without fear.

Governance should enable momentum, not freeze it. If governance is unclear, employees guess. If governance is too heavy, experimentation moves underground. The right model makes safe experimentation visible, keeps learning in the open, applies risk tiers to different use cases, and creates a cross-functional decision rhythm for tools, data, pilots, policy refinement, and adoption. The goal is safe momentum, not static control.

The practical message is simple: AI use policy, safe-use training, and governance are not administrative side work. They are the operating conditions that allow AI activity to become responsible business value.

Contents

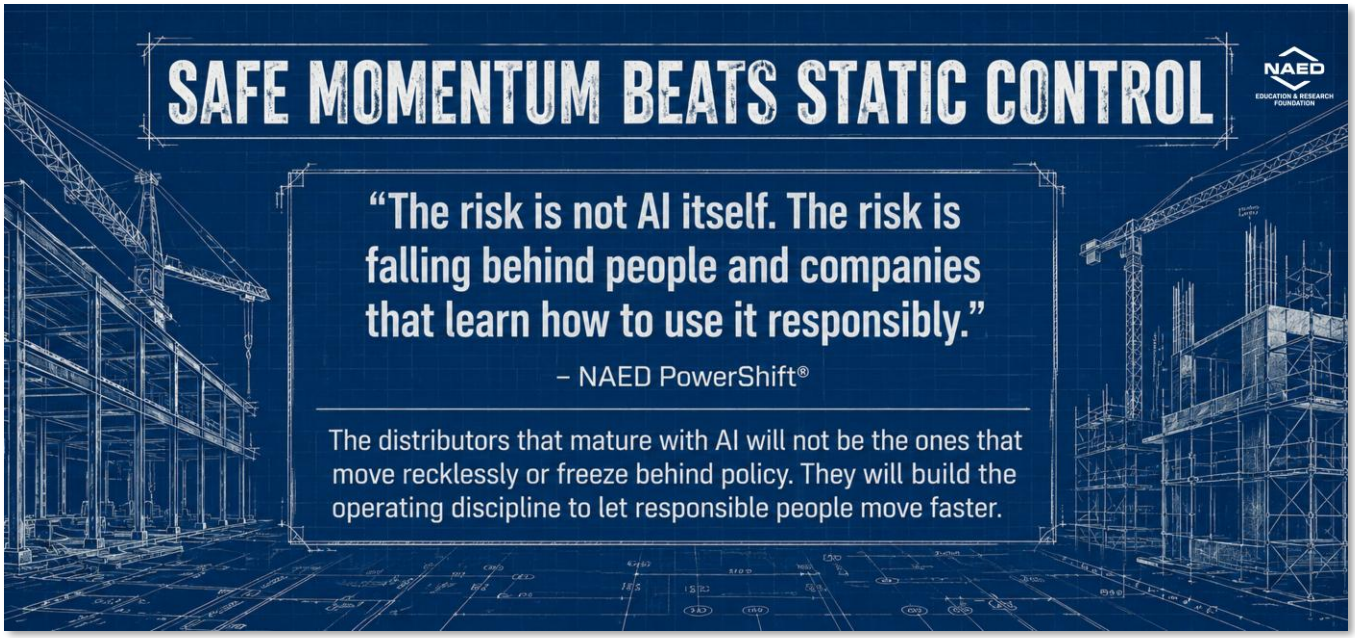

AI Use Policy as an Operating Tool

An AI use policy should not be written as if the main audience is the legal department. Legal, compliance, IT, security, HR, and executive leadership all have a role in shaping the policy, but the policy succeeds or fails with employees and managers. If a counter salesperson, buyer, warehouse supervisor, branch manager, HR generalist, marketing specialist, or credit analyst cannot understand what the policy means in their work, the policy is not finished.

The strongest AI use policy starts with the operating reality of the distributor. Distributor work is fast, practical, and relationship-driven. Customers ask for answers before all information is perfectly organized. Sales teams interpret order status, availability, substitutions, lead times, and customer expectations. Buyers compare supplier updates and shortages. Operations teams manage exceptions. Warehouse teams execute under time pressure. Managers translate messy inputs into decisions. AI can assist that work, but it can also create confident output from weak assumptions. The policy has to help people use AI without weakening trust.

The policy should define the difference between approved tools and approved uses. This distinction matters. A company may approve an enterprise AI tool, but that does not mean every kind of data belongs in it. Tool approval is only one layer. Data boundaries, use-case risk, review expectations, ownership, and escalation are separate questions. Employees need to understand that an approved tool does not automatically make a high-risk use case safe.

A practical policy should make data boundaries concrete. Do not rely on broad phrases like confidential information or sensitive data without examples. In the electrical distribution environment, employees need explicit guidance around customer PII, employee information, payroll data, credit information, pricing logic, margin strategy, customer contracts, supplier agreements, project bids, confidential financials, cyber/security information, proprietary product content, and non-public operational data. The policy should tell people what is never allowed, what may be allowed only in approved tools, and what requires review before use.

Human review must be a core policy requirement. AI can draft, summarize, organize, compare, and suggest. It should not own customer commitments, HR decisions, legal interpretation, pricing recommendations, safety guidance, credit decisions, margin-impacting actions, or supplier commitments. The human owner must remain responsible for facts, source quality, tone, assumptions, business judgment, and final action. That review standard is not a technical detail. It is how the distributor protects customer trust and operating quality.

The policy also needs an escalation path. Employees should not have to guess what to do when a use case feels useful but unclear. The safe answer should not be hidden in a document, buried in Teams, or dependent on knowing the right person to ask. A practical escalation path may begin with a manager, AI champion, IT/security contact, data owner, legal/HR owner, or AI council, depending on the issue. The key is that asking should be normal and encouraged. Hidden uncertainty is more dangerous than visible questions.

A practical AI use policy should answer these questions in plain language:

Policy Area	Question the Policy Must Answer	Why It Matters
Tools	Which AI tools are approved, prohibited, or pending review?	Prevents tool sprawl and unmanaged usage.
Data	What data can be used, restricted, or never entered?	Protects customer, employee, supplier, financial, and proprietary information.
Use cases	Which uses are low, medium, or high risk?	Matches controls to business impact instead of treating all AI use the same.
Human review	What outputs must be reviewed before use?	Keeps people accountable for commitments, decisions, and quality.
Escalation	Who should employees ask when unsure?	Removes ambiguity before it turns into shadow AI or unsafe workarounds.

Training Safe AI Use Like Cybersecurity

Safe AI usage training should be designed like cybersecurity training because both depend on daily behavior. A cybersecurity policy does not protect the business if employees cannot recognize risky behavior, respond to uncertainty, or report a mistake. AI is similar. The risk often appears in small decisions made under time pressure: pasting the wrong data, accepting the first answer, forwarding a weak draft, relying on an unsupported summary, or using an unapproved tool because it is convenient.

The old training model will not be enough. A one-time launch event, static policy acknowledgement, or generic software training video may create awareness, but it will not create durable behavior. Distributors need repeated, scenario-based, role-specific training that helps employees recognize the safe path inside their own work. Counter sales needs different examples than HR. Purchasing needs different examples than marketing. Managers need different guidance than frontline users. Executives need different decision standards than individual contributors.

Training should begin with the tool boundary, but it cannot stop there. Employees should know which tools are approved, what data restrictions apply, and why personal or public tools may create exposure. They also need to understand that AI output is not self-validating. The fact that an answer is well written does not make it accurate. The fact that a summary is confident does not make it complete. The fact that a draft sounds professional does not make it safe to send.

The strongest safe-use training teaches employees to critique the output. Before using AI output, employees should ask what assumptions are being made, what information may be missing, what source material the answer relies on, what facts need verification, what tone or commitment risk exists, and whether the output is appropriate for the audience. That second-pass critique is one of the most practical habits a distributor can teach. It converts AI from a one-click answer machine into a review-supported work practice.

Training also has to address work slop. AI can help good employees produce more output faster, but more output is not the same as better work. Work slop appears when polished drafts, summaries, reports, or slide decks create the illusion of progress without supporting a decision, improving a workflow, reducing rework, or serving a customer. Safe AI training should teach employees to define the job before generating output:

- What decision are we supporting?
- Who is the audience?
- What would make this useful?
- What must be verified?
- What should be omitted?

Managers are the reinforcement layer. If managers do not understand safe AI use, adoption and behavior will be uneven. Some employees will overuse AI, some will avoid it, and some will use it quietly without guidance. Managers need short discussion guides, examples of appropriate and inappropriate use, escalation paths, and permission to surface confusion. Training should equip managers to coach judgment, not merely enforce rules.

Training should make five behaviors automatic:

- Use approved tools for company work and know when tool approval does not mean data approval.
- Protect restricted data, including customer, employee, supplier, pricing, contract, financial, and proprietary information.
- Define the task, audience, decision, and review standard before generating AI output.
- Critique AI output for assumptions, missing context, unsupported claims, and risky commitments.
- Escalate uncertainty quickly, visibly, and without fear of being punished for asking.

Governance That Enables Momentum

Governance is often misunderstood. Some employees hear governance and assume the company is slowing everything down. Some leaders hear governance and think the problem is solved once a policy is published. Both views are incomplete. Governance is not a stop sign and it is not a static document. Good governance is the operating discipline that makes safe progress repeatable.

The best governance model is business-led and IT-partnered. The business owns the pain, workflow, value, change management, data cleanup priority, and adoption. IT and data leaders manage infrastructure, security, access, architecture, integration reality, monitoring, and technical controls. Neither side can do this alone. If the business throws vague AI ideas to IT, value will be weak. If IT builds guardrails without business context, governance will feel disconnected from the work. The model needs both engines.

Risk tiering is central. A distributor should not govern every AI use case the same way. Personal productivity use, such as brainstorming or summarizing non-sensitive notes, may require basic review and approved tool use. Workflow support, such as drafting customer follow-ups or summarizing supplier communication, may require approved sources, manager review, and clear human approval. Business-impacting decisions, such as pricing, credit, HR, legal, safety, customer commitments, or automated system actions, require stronger ownership, review, controls, and evidence.

Governance also needs an escalation funnel. The frontline employee should not carry the full burden of interpreting every policy question. The manager should be able to answer common questions. AI champions can help with examples and peer learning. A cross-functional governance council can resolve recurring ambiguity, update policy, identify risks, and decide when a use case needs deeper review. This funnel removes ambiguity without forcing every question into executive leadership or IT.

A practical governance rhythm should evolve over time. In early maturity, an AI Champions Group may be enough to make learning visible, share examples, capture prompt patterns, and surface questions. As business value and risk increase, the organization may need a more formal AI council that reviews use-case intake, data readiness, tool requests, pilot outcomes, risk tiers, training needs, and adoption challenges. The point is not to create a committee for its own sake. The point is to create a trusted decision rhythm.

Governance should also support the role of consultants clearly. Consultants can help benchmark what peers are doing, facilitate policy design, build training, map use-case risk, develop governance structures, and accelerate internal clarity. They should not own the business case, replace leadership judgment, or become the permanent decision body. A good consultant leaves the distributor with stronger internal capability and a clearer operating rhythm.

Governance should make these decisions visible:

Governance Decision	Question	Purpose
Tool approval	Which tools are approved, prohibited, or pending review?	Prevents unmanaged tool sprawl.
Data boundaries	Which data types are safe, restricted, or prohibited?	Protects customers, employees, suppliers, and proprietary business knowledge.
Use-case risk	Which uses require manager review, functional approval, or formal governance?	Protects business-impacting decisions.
Human review	Who approves output before action or external use?	Prevents diffusion of responsibility.
Ownership	Who owns the policy, the use case, the data source, and the adoption path?	Keeps AI from becoming everyone's job and no one's accountability.
Measurement	How will we know whether AI is improving the work?	Connects activity to operational value.

Conclusion: Safe Momentum Beats Static Control

AI use policy, safe-use training, and governance should not be treated as defensive administration. They are part of the operating model required to turn AI curiosity into safe, practical value. A distributor cannot mature if employees are guessing about tools, managers are unsure what to reinforce, IT is expected to own business transformation alone, or leadership assumes that a policy document equals adoption.

Static control creates false confidence. It looks responsible, but it often fails where the work actually happens. Employees still face daily decisions about what data to use, what output to trust, what customer language to send, what analysis to rely on, and when to ask for help. A static policy cannot handle that alone. Safe momentum comes from clear guardrails, repetitive scenario-based training, manager reinforcement, visible experimentation, risk-tiered governance, and a rhythm for improving the system as the organization learns.

The strongest distributors will not be the ones with the longest AI policies. They will be the ones whose people understand the safe path, whose managers reinforce good behavior, whose leaders connect AI to real friction, and whose governance model keeps learning visible.

That is the practical standard for NAED PowerShift®: move fast enough to learn, safely enough to protect trust, and deliberately enough to turn activity into measurable business value.

Appendix A: AI Use Policy Starter Framework

This starter framework is not a legal template. It is a practical structure leaders can use to shape a policy that is clear enough for employees and managers to apply. Legal, HR, IT, security, and compliance review should be added before formal adoption.

Policy Element	Starter Guidance	Practical Note
Purpose and principles	State that AI is encouraged when used responsibly to improve work, support judgment, and protect trust.	Keep the tone enabling, not fear-based.
Approved tools	List approved tools, prohibited tools, and the process for requesting new AI-enabled tools.	Update frequently. Tool lists age quickly.
Data boundaries	Define data that is allowed, restricted, or prohibited. Include distributor-specific examples.	Name customer, employee, supplier, pricing, contract, financial, and operational examples.
Risk tiers	Separate low-risk productivity, medium-risk workflow support, and high-risk business-impacting use.	Controls should match risk.
Human review	Define outputs that require human approval before use.	Customer-facing, pricing, HR, legal, safety, credit, and operational commitments need review.
Accountability	Clarify that AI assists work, but people own decisions, quality, and commitments.	Avoid diffusion of responsibility.
Escalation	Provide a clear path for questions, uncertainty, mistakes, and new use cases.	Asking should be expected.
Review cadence	Set a review rhythm for policy updates, training refreshes, and governance decisions.	AI policy should evolve with use.

Appendix B: AI User Training Starter Curriculum

This curriculum should be delivered in short sessions, reinforced by managers, and refreshed regularly. The goal is behavior change, not policy awareness alone.

Training Module	Purpose	Primary Audience
Module 1: Why Safe AI Use Matters	Explain opportunity, risk, and the distributor-specific reason safe use matters.	All employees
Module 2: Approved Tools and Data Boundaries	Show which tools are approved and which data types require protection.	All employees
Module 3: Prompting as Work Definition	Teach users to define purpose, audience, constraints, and review expectations.	All AI users
Module 4: Human Review and Second-Pass Critique	Teach users to check assumptions, missing data, overstatements, and unsupported claims.	All AI users
Module 5: Role-Based Scenarios	Use practical examples for sales, purchasing, operations, warehousing, HR, finance, marketing, and managers.	Functional teams
Module 6: Work Slop and Quality Standards	Show how polished AI output can still be low-value, incomplete, or risky.	Managers and content-heavy roles
Module 7: Escalation and Reporting	Teach employees how to ask questions and report mistakes or uncertainty.	All employees
Module 8: Manager Reinforcement	Equip managers with team huddle questions, examples, and coaching language.	Managers

Appendix C: Governance Questions for Leadership

Leadership should use these questions to assess whether the organization has enough governance to support safe momentum. A strong answer does not require bureaucracy. It requires clarity, ownership, and a decision rhythm.

Tools and access

- Which AI tools are approved for company work?
- Which tools are prohibited or require review?
- Who approves AI-enabled features in existing software?
- How do employees find the current approved-tool list?

Data and information boundaries

- What information can never be placed into an AI tool?
- What information can be used only in approved enterprise tools?
- Who owns decisions about customer data, employee data, pricing, contracts, supplier information, and financial data?
- How are data boundaries explained through practical examples?

Human review and accountability

- Which outputs require review before use?
- Who owns customer-facing, pricing, HR, legal, safety, credit, and operational commitments?
- How should reviewers check facts, assumptions, tone, source quality, and business impact?
- How will the company prevent AI output from creating work slop?

Governance and escalation

- Who owns AI governance today?
- Does the organization need an AI Champions Group, AI council, or both?
- What questions should managers answer versus champions versus IT/security versus executives?
- How will recurring ambiguity become policy or training updates?

Measurement and improvement

- How will leaders know whether AI use is improving work?
- What early signals matter: time saved, rework reduced, response speed, quality, adoption, fewer escalations, or better consistency?
- How will pilots or high-value use cases be reviewed?
- What will trigger a stop, refine, or scale decision?

Consultant role

- Where does outside help accelerate clarity or capability?
- What should remain owned internally by the business, IT, HR, legal, and leadership?
- How will consultants transfer knowledge rather than create dependency?
- What final artifacts should consultants leave behind: policy draft, training material, governance charter, use-case triage model, or roadmap?

Appendix D: Infographic