
A FRAMEWORK FOR EVERY
ORGANIZATION

The AI Policy *Standard*

Seven components every organization needs — and
the plain-English questions leadership has to answer.

Chuck Gallagher
ETHICS + AI KEYNOTE SPEAKER

FIELD GUIDE
№ 02

THE WATER YOU'RE ALREADY IN

Every organization already uses AI. Very few have written down the rules.

Copilot rewrites your emails. Gemini answers your searches. ChatGPT drafts your proposals. The technology arrived without an announcement — and the policy that governs it hasn't been written yet.

The question is no longer whether the organization uses artificial intelligence. It does. Every knowledge worker with a Microsoft license, a Google Workspace account, or an internet browser is using AI several times a day.

The question is what happens when that work goes out the door. When the proposal reaches the customer with a number the AI made up. When the meeting summary contains a commitment nobody remembers making. When the handbook posted on the company intranet was drafted by a tool that had no idea what state the company operates in.

Those things are happening at organizations right now. Some get caught. Many don't. The ones that don't are the reason the phone rings.

This field guide is not a legal document. It is the plain-English framework I put on the table in the first working session with every client who asks a version of the same question — *Chuck, what should actually be in our AI policy?*

Seven components. Every organization needs them. What follows.

THE COST OF THE MISSING PAGE

What a missing policy actually costs.

A distribution company outside Charlotte sends a proposal to a longtime customer. Someone on the team used AI to help draft it. The proposal contains a number. Nobody verifies the number before it goes out under the owner's signature. The customer catches it. A phone call follows. A hard one.

Fortunately, the customer was willing to hear her out. Had the customer held firm, the outcome would have been very different — and costly to the firm.

That is the specific moment a policy is designed to prevent. Not through more technology. Through a written expectation that human eyes review AI-assisted work of a certain kind before it leaves the building.

Most organizations have not written that expectation down. Their teams are using AI dozens of times a day, without guidance, without a review step, without a defined line between what the tool can produce and what a human has to approve. The gap is not a policy problem waiting to happen. It is a policy problem happening right now, quietly, at every level of the organization.

The seven components on the pages that follow are the framework that closes that gap. They are not a compliance manual. They are the shape of the conversation leadership has to have — with counsel, with the board, and with itself — before the next avoidable mistake goes out under the company's name.

The Seven Components.

The framework — first four.

01 Scope

What AI tools does the policy cover?

Not just the tools someone signed a contract for. Copilot in Outlook and Word. Gemini inside Google. AI features baked into Zoom, Teams, Salesforce, the CRM, the accounting software. Almost every enterprise application now contains an AI feature the vendor turned on by default. The policy names them all.

The mistake: Policies that only cover ChatGPT while Copilot rewrites every email in the building.

02 Approved Use Cases

What is the team allowed to use AI for?

Drafting emails. Summarizing meetings. First-draft copy for internal documents. Research and background reading. Named in a sentence a new hire can read on day one and understand. Approved use cases give the team permission to move quickly on the work AI is genuinely good at.

The mistake: Silence — which employees read as either "everything is allowed" or "nothing is allowed," neither of which is true.

03 Prohibited Use Cases

What is AI not allowed to touch?

Confidential customer data. Employee records. Financial figures that haven't been verified. Legal or medical decisions. Anything that ends up in front of a regulator. Prohibited use cases are the guardrails, and they need to be as specific as the approved list.

The mistake: Prohibitions written in language so broad ("do not misuse AI") that no employee can tell whether their next action crosses the line.

04 The Review Step

Who checks AI-assisted work before it goes out the door?

The section that would have caught the wrong number in the Charlotte proposal. Names the second set of eyes for each category of AI-assisted work — customer-facing documents, financial figures, contracts, public statements. Names how the review happens. Names when it can be waived, and by whom.

The mistake: No policy at all, followed by an argument about who should have caught it.

The Seven Components.

The framework — final three.

05 Disclosure

When does the organization tell others that AI was used?

Most of the time, nothing needs to be said — nobody discloses spell-check. Sometimes the answer changes. Contracts. Legal correspondence. Advisory work. Medical or financial guidance. Public statements. Any work where the audience is entitled to know a human made the judgment call. The policy names those moments in advance so the team doesn't have to decide in the moment.

The mistake: Treating every AI-assisted document as if it needs a disclosure — or, more commonly, treating none of them as if they do.

06 Data Handling

What information is not allowed into a public AI tool?

Customer names. Financial data. Employee personal information. Anything under NDA. Anything covered by HIPAA, FERPA, GLBA, or a state privacy statute. The paragraph that keeps the organization out of a news story it doesn't want to be in — and out of a regulatory action it can't afford.

The mistake: Assuming enterprise contracts with vendors solve the problem when the employee is using the consumer version on a phone.

07 Ownership and Accountability

Who owns the outcome when AI-assisted work goes wrong?

The employee who used the tool. The manager who approved the output. The department head who signed off on the process. Named in advance, in writing, so nobody has to figure it out at 4:47 on a Friday afternoon after a customer email has already gone out. Ownership on paper before ownership becomes a legal question.

The mistake: A policy that names the tools but not the people, and leaves ownership to whoever the finger points at first.

FROM FRAMEWORK TO PRACTICE

Five steps from framework to practice.

The framework is portable. The policy is not. Every organization writes its own version — because industry, size, regulatory posture, and customer base change the language on the page. What does not change is the sequence of work required to get from an idea into a document that leadership actually enforces.

1 Convene the working group.

Legal, HR, IT, operations, and one senior leader with authority to make decisions. Five people. Two hours. Not a committee. A working group with a deadline.

2 Audit what's already in use.

Before you write a policy, list every AI tool employees are already using — including the ones running in the background of software the organization already owns. The list is almost always longer than leadership expects.

3 Draft the seven sections plainly.

One page per component. Written in the language a new employee could understand on day one. No compliance jargon. No hedges. The document is a working tool, not a shield.

4 Review with counsel — and then trim.

Counsel will add. That is their job. Leadership's job is to trim it back to something the team will actually read. A policy nobody reads is a policy that isn't in force.

5 Train, publish, and set a review date.

Every employee sees the policy in the first thirty days. The document has a review date on the last page — no more than twelve months out. AI is moving too quickly to leave a policy on a shelf.

TEN QUESTIONS FOR LEADERSHIP

The self-audit.

Before the working group meets, work through the ten questions below. The answers become the starting outline for the policy. Any question that produces the answer "I'm not sure" is a section that needs written attention this quarter.

BOARD-LEVEL DIAGNOSTIC

Ten questions leadership has to answer.

Answer honestly. The gaps are the policy.

- 01** Do we have a written list of the AI tools employees are permitted to use — including the ones already built into software we own?
- 02** Can a new hire describe, in one sentence, what they are allowed to use AI for on day one of the job?
- 03** Can that same new hire describe, in one sentence, what they are not allowed to use AI for?
- 04** Is there a review step — a named human — for AI-assisted work that touches customers, contracts, money, or regulated information?
- 05** Do we know, as leadership, when we tell a customer, member, or partner that AI was involved in the work we sent them?
- 06** Is there a written list of the information categories that are not permitted to be entered into a public AI tool?
- 07** If an AI-assisted mistake leaves the building tomorrow morning, do we know — in advance — who owns the outcome?
- 08** Has counsel reviewed our current AI practices in the last twelve months?
- 09** Has every employee been trained on the policy in the last twelve months?
- 10** Do we have a review date on the calendar for the next version of this policy?

ABOUT THE AUTHOR

The conversation your leadership hasn't had yet.

Chuck Gallagher is an ethics and AI keynote speaker, consultant, and author of six books including *AI Real Time*, *Every Choice Has a Consequence*, and *The Technology Advantage*. He works with associations, boards, and leadership teams on the practical ethics of adopting AI without letting the mistakes go out the door under the company's name.

This field guide is one in a series. It is designed to be shared — sent to a board chair, forwarded to a general counsel, distributed to the members of an association that needs to get in front of the AI conversation before somebody else does.

SPEAKING

Book a keynote

Association meetings, corporate leadership retreats, board briefings, and conference plenaries on AI ethics and organizational adoption.

CONSULTING

Start with a call

Working sessions to help leadership teams write, refine, and roll out an AI policy that fits the organization. The first fifteen minutes are free.