

THE BUSINESS LEADERS SERIES

AI for Business Leaders

GETTING REAL VALUE FROM AI WITHOUT GETTING
BURNED BY IT

A PLAIN-LANGUAGE EXECUTIVE GUIDE

*AI is a powerful assistant
and a poor authority.*

Build on that one sentence and you keep the value without most of the harm.



THIS GUIDE · 4 OF 6 IN THE SERIES

J.W. WALKER · [JEREMYWALKER.US/RESOURCES](https://jeremywalker.us/resources)

AI is a business decision, not just a technology one. Every leader is being told they must adopt AI immediately or be left behind. Much of that pressure comes from people selling AI. The real task is quieter and harder: figure out where AI genuinely helps your business, adopt it in a way you can control, and govern the risk it introduces. This guide is written to help you do all three, in plain terms, without the hype and without the fear.

TWO IDEAS TO HOLD AT ONCE

AI rewards leaders who can hold two ideas at the same time. Drop either one and you get into trouble.

Adoption and governance are the same conversation. The same qualities that make AI powerful, its speed, its reach, its ability to act on data, are exactly what make it risky when it is wrong or misused. You cannot adopt it well without governing it, and governing it is not a reason to avoid it. Leaders who chase adoption without governance get burned. Leaders who fear the risk and avoid it get left behind. The job is both.

A decision you cannot explain is a decision you cannot defend. When AI influences something that matters, a loan, a hire, a fraud flag, a member communication, you have to be able to explain how the decision was made, to a regulator, to a customer, and to yourself. If the answer is “the model decided,” you do not have a defensible position. This single principle governs most of what follows.

WHAT'S IN THIS GUIDE

1. What AI Actually Is	The plain-language version, without the jargon or the hype
2. Where AI Creates Real Value	The uses that pay off, what they cost, and how to spot them
3. AI and Your People	What AI means for your workforce, and the honest leadership position
4. Where AI Is Oversold	The hype, the limits, and the failure modes to watch for
5. Buying AI: Build, Buy, or Embedded	How to evaluate AI tools and what "AI-powered" really means
6. The Risks You Now Own	Shadow AI, data leakage, output ownership, and third-party AI
7. AI as a Threat	Deepfakes, AI-enabled fraud, and attacks against your business
8. Governing AI in Practice	A policy starter, the guardrails, and who owns it
9. AI and the Regulators	What financial institutions and regulated businesses must know
10. Questions Leaders Should Ask	A practical checklist before and during adoption
11. Getting Started Without Getting Burned	A measured path to real value

1. What AI Actually Is

You do not need to understand how AI works internally any more than you need to understand how an engine works to run a trucking company. But you do need an accurate mental model, because many bad AI decisions trace back to a wrong one.

Strip away the marketing and much of modern AI, especially the generative kind everyone is talking about, works by learning patterns from enormous amounts of data and using those patterns to predict or generate outputs. A large language model, the technology behind tools like ChatGPT, predicts the next likely piece of text in a sequence. That is a narrow description of something genuinely powerful, but it is the honest one, and it explains both what AI is remarkable at and where it fails.

The Vocabulary, Briefly

Artificial Intelligence	The broad field of software that performs tasks we associate with human intelligence. An umbrella term, often used loosely.
Machine Learning	Systems that learn patterns from data rather than following rules a person wrote. The engine under most modern AI.
Training Data	The information used to teach an AI system the patterns it later relies on. Its quality, completeness, relevance, and bias can materially shape the results.
Predictive AI	The category of AI that analyzes patterns to classify, score, forecast, or recommend. Common examples: fraud detection, risk scoring, forecasting.
Generative AI	The category of AI that produces new content: text, images, code, audio. The category behind the current wave of tools.
Agentic AI / AI Agent	AI that can plan, use tools, and take multi-step actions, not just produce an answer. Wrong output can become a wrong action.
Model	The trained system itself. When people say "the model decided," they mean this. It is software, not a mind.
Prompt	The instruction or question you give an AI tool. The quality of the prompt shapes the quality of the output.

Human in the Loop	A person who reviews, approves, or can override AI output before it becomes consequential. It only works as a control if that person has the expertise and authority to challenge the AI.
Hallucination	When an AI produces confident, fluent output that is simply wrong. A structural feature, not an occasional bug.
Large Language Model (LLM)	A generative AI trained on vast text to predict and produce language. Powers chatbots and writing assistants.

One Thing to Remember

AI predicts plausible output. It does not know things the way a person does, and it has no built-in sense of whether what it produced is true. It can write a flawless-sounding paragraph that is entirely fabricated, cite a regulation that does not exist, or state a confident number that is wrong. This is not a failure mode leaders should assume has been eliminated simply because models keep improving; it reflects how the technology works. That single fact should shape every decision about where you let AI operate and how much you check its work.

The Core Principle

AI is a powerful assistant and a poor authority. Use it to draft, summarize, accelerate, and suggest. Do not use it as the final word on anything that matters without a human who is accountable for the result.

2. Where AI Creates Real Value

AI is neither magic nor a gimmick. It creates genuine value in specific, identifiable situations, and understanding the pattern helps you find the opportunities in your own business rather than chasing whatever a vendor is promoting.

The common thread across valuable AI uses is this: tasks that involve processing language or patterns at volume, where a knowledgeable human reviews the result. AI can handle much of the first-pass work quickly; a knowledgeable person provides the judgment, correction, and accountability before the result matters.

Where It Genuinely Helps

- **Drafting and summarizing.** First drafts of communications, summaries of long documents, meeting notes, and reports. The human edits and owns the final version, but starts from something rather than a blank page.
- **Customer and member service support.** Answering routine questions, routing inquiries, and helping staff find information faster. Best when it assists a human rather than replacing them on anything sensitive.
- **Analyzing patterns in data.** Fraud detection, anomaly spotting, and surfacing trends in large datasets. AI is genuinely strong here, which is also why attackers use it.
- **Coding and technical work.** Helping developers write, review, and debug code faster. One of the more established productivity uses, with human review still essential.
- **Search and knowledge access.** Helping employees find answers buried in policies, manuals, and internal documents far faster than keyword search.

How to Spot a Good Use Case

A strong AI opportunity usually has three features. The task involves language, patterns, or large volumes of information. A knowledgeable person can quickly check whether the output is good. And a wrong answer is inconvenient rather than catastrophic, or is caught before it causes harm. When all three are present, AI tends to pay off. When the task demands guaranteed accuracy, accountability with no human in the loop, or a decision that must be explained precisely, proceed with far more caution.

The best early AI wins are usually internal and low-stakes: helping your own people work faster on tasks they already understand well enough to catch mistakes. Start there, not with a customer-facing system making consequential decisions.

What It Actually Costs

AI value is not free, and the sticker price is rarely the real cost. Most business AI is sold on a subscription, either per seat or by usage, and usage-based pricing can climb unpredictably as adoption grows. But the cost that catches organizations off guard is not the license. It is the human review that responsible AI use requires. If a tool drafts a document in seconds but a person still has to read, correct, and own it, you have shifted the work rather than eliminated it, and that time is a real cost. The genuine savings come when AI reliably handles volume that a person then spot-checks rather than fully reworks.

The honest way to evaluate return is not “how impressive is the demo” but “how much real time or error does this remove, net of the review it requires, on our actual work.” A tool that saves an hour of drafting but adds forty minutes of checking is a very different investment than the vendor’s slide suggests. Pilot on real tasks, measure the net, and be skeptical of savings that only appear in the sales deck.

3. AI and Your People

The first question most boards and executives ask about AI is not technical. It is human: what does this mean for our people? It deserves a direct answer, because how you handle it shapes both the value you get from AI and the trust you keep inside your organization.

For most organizations, the strongest near-term use of AI is augmentation first: removing low-value work and increasing the output of capable people. AI is genuinely good at pieces of jobs, drafting, summarizing, searching, first-pass analysis, and genuinely poor at owning a job end to end, because it cannot be accountable. But augmentation is not a promise that staffing never changes. AI can reduce future hiring, consolidate work, reshape roles, and in some cases eliminate them. The honest leadership position is to lead with augmentation while being straight that job design and staffing may shift over time, rather than promising outcomes you cannot guarantee.

What Changes for Your Workforce

The work shifts toward judgment. As AI handles more of the drafting and processing, the human value moves to reviewing, deciding, and owning outcomes. That is a real change in what you are hiring and training for, and it rewards people who can evaluate AI output critically rather than produce the first draft themselves.

AI literacy becomes a baseline skill. Your people do not need to become data scientists, but they do need to understand what AI is reliable for, where it fails, and what data they may and may not enter into specific AI tools. That is a training obligation, and it is also your first line of defense against the data-leakage risk covered later in this guide.

Trust and transparency matter. Staff who fear AI is coming for their jobs will either resist it or use it in the shadows. Leaders who are honest about how AI will and will not be used, and who involve employees in finding good uses, get better adoption and fewer nasty surprises.

The Leadership Position

The most defensible stance for most organizations: AI is augmentation first, removing low-value work and raising what we expect of capable people, and it may still change hiring, staffing, and job design over time. Say both parts honestly. Promising that nothing will change is not kinder; it is less credible, and it erodes the trust you need for people to bring their AI use into the open.

4. Where AI Is Oversold

Alongside genuine AI use cases are many that are exaggerated, premature, or simply wrong. Knowing where AI is oversold protects you from wasting money, damaging trust, and making decisions you cannot defend.

The Failure Modes to Watch

Confident wrongness. AI produces fluent, authoritative output whether or not it is correct. It does not signal uncertainty the way a careful person does. The more plausible the output looks, the more dangerous a wrong answer becomes, because no one thinks to check it.

The accountability gap. When AI makes or shapes a decision, who is responsible when it is wrong? “The system did it” is not an answer a regulator, a court, or a wronged customer will accept. If you cannot name the human accountable for an AI-influenced outcome, you have a governance problem, not a technology one.

The explainability problem. Many AI systems cannot fully explain why they produced a particular result. For consequential decisions, especially in regulated areas like lending or hiring, “we cannot explain how the model reached this” is a serious liability, not a technical footnote.

Bias at scale. AI learns from data, and if that data reflects historical bias, the AI can reproduce or amplify it, quickly and consistently, though with careful design, data, and validation it can also reduce some forms of human inconsistency. In lending, hiring, and pricing, poorly controlled bias can create serious ethical problems and legal exposure under fair-lending and anti-discrimination law.

The demo-to-production gap. AI tools demo beautifully on curated examples and behave very differently on the messy reality of your actual data and edge cases. An impressive demonstration is not evidence the tool will work reliably in your environment.

The Hype Vocabulary

When a vendor says a product is “AI-powered,” that phrase can mean anything from a genuinely sophisticated system to a thin wrapper around someone else’s model to a simple rules engine with a marketing label. The term itself tells you almost nothing. The useful questions are what the AI actually does, what it was trained on, how errors are handled, and who is accountable when it is wrong. A vendor who cannot answer those plainly is selling the label, not the capability.

A Useful Filter

When you hear an AI claim, ask: what specifically does it do, how do we know when it is wrong, and who is accountable when it is? Vague answers to those three questions are a clear signal that a product is more hype than substance.

5. Buying AI: Build, Buy, or Embedded

Most organizations will not build their own AI. They will buy it, or more often, it will arrive embedded inside software they already use. Understanding the three paths helps you make deliberate decisions rather than accumulating AI by accident.

The Three Paths

Build. Developing custom AI in-house or with a partner. Expensive, slow, and requiring genuine expertise, this is rarely the right first step and is only justified when AI is core to your competitive advantage and no existing tool fits. Most organizations should not start here.

Buy. Purchasing a dedicated AI tool or platform for a specific purpose. The most common deliberate path. Your leverage is in evaluation and contract terms: what the tool does, how your data is handled, and what happens when it is wrong.

Embedded. AI features arriving inside tools you already have: your core platform, your office software, your customer service system. Many organizations encounter AI this way without

deliberately choosing it, which can introduce ungoverned use because the capability simply appeared in an update.

Questions Before You Buy

- What specifically does the AI do, and what problem of ours does it solve?
- What data does it use, where does that data go, and is any of it used to train the vendor's models?
- How are errors detected and handled? What happens when the AI is wrong?
- Can the vendor explain how the system reaches its outputs, to the level a regulator would expect?
- What are the vendor's security practices, and how are they subject to our third-party risk review?
- If AI influences a consequential decision, what human review and override exists?
- What is our exit path if the tool underperforms or the vendor changes terms?
- How will we be notified if the underlying model, material functionality, data practices, or subprocessors change?

The buying question that matters most for a regulated or data-sensitive business: does our data leave our control, and is it used to train the vendor's models? If you do not know the answer, you are not ready to buy.

6. The Risks You Now Own

The moment AI enters your organization, whether you chose it or it arrived embedded, you own a set of new risks. Most of them are invisible until something goes wrong, which is precisely why they need deliberate attention.

Shadow AI

Your employees are very likely using AI tools already, whether sanctioned or not. They use consumer AI assistants to draft emails, summarize documents, analyze spreadsheets, and answer work questions. This is shadow AI: use that is invisible to leadership and therefore ungoverned. The risk is not that employees use AI. It is that no one knows what data is going into which tools, or what happens to it once it does.

Data Leakage

When an employee pastes customer records, financial data, internal strategy, or sensitive documents into a public AI tool, that information may, depending on the provider, service tier, and configuration, be retained, used to improve the provider's models, or exposed in ways that violate privacy obligations, contracts, or regulation. Data handling varies substantially between a free consumer tool and a properly contracted enterprise service, which is exactly why the tool and the use both have to be approved rather than assumed safe. The employee is usually just trying to work faster and has no idea they created exposure, and once sensitive data has been entered into an external system, you often cannot get it back. This is a common and frequently underestimated AI risk.

Ownership of What AI Produces

When AI generates content your business uses, a marketing draft, code, an image, an analysis, two ownership questions follow, and both are unsettled enough to warrant caution. First, can you actually own or protect what the AI produced? Purely AI-generated material generally is not protected by copyright without sufficient human authorship. Human-AI works may still qualify to the extent that a person contributes protectable expression, selection, arrangement, or modification. [8] This matters if the output is something you intend to own exclusively. Second, and more immediately, does using the output expose you to someone else's claim, if it incorporates or is substantially similar to protected material? For most everyday uses this is low risk, but for anything you will publish, sell, or build a product around, it is worth a deliberate check rather than an assumption. Read your AI tools' terms on who owns outputs and what indemnity, if any, the vendor provides.

The Ownership Question

Before you build something commercial on AI-generated output, ask two things: can we own and protect this, and could using it infringe on someone else's rights? For routine internal drafting the answer rarely matters. For anything you will publish or sell, get a clear read on the tool's terms before you rely on it.

Model and Decision Risk

When AI influences decisions, its errors and biases become your errors and biases, at scale and speed. A biased lending model does not make one unfair decision; it makes thousands

consistently. An AI that subtly drifts as conditions change can degrade without anyone noticing until the pattern of bad outcomes is large enough to see. Any AI that touches a consequential decision needs monitoring, testing, and a human accountable for the results.

Third-Party and Embedded AI Risk

Much of your AI exposure will come through vendors: AI features in your existing platforms, or AI-driven services you rely on. You inherit the risk of AI you did not build and may not fully understand. If a vendor's AI mishandles your data or produces a discriminatory outcome using your information, the accountability does not stop at the vendor's door. AI vendors belong squarely inside your third-party risk program, evaluated with the same rigor as any other provider with access to your data.

The Underestimated Risk

Ask your team a direct question: what AI tools are our employees using right now, and what data has gone into them? In many organizations, leadership does not know the full answer. That gap is where the risk lives, and closing it is the first governance task, not the last.

7. AI as a Threat

AI does not only create internal risk. It hands new capabilities to the people trying to defraud, deceive, and attack your business. These threats are already in use against financial institutions and businesses of every size, and awareness is the first line of defense.

Deepfake Voice and Video

AI can now clone a person's voice from a small sample of audio, often taken from a voicemail, a recorded call, a podcast, or a social media clip. Attackers use cloned voices to impersonate executives authorizing urgent transfers, board members making unusual requests, or customers bypassing phone-based verification. Deepfake video is advancing quickly along the same path. The federal financial-crimes network has warned specifically about deepfake media being used to defeat identity verification and commit fraud against financial institutions. [7]

Leadership Awareness

If you have a public profile, usable samples of your voice may already be publicly available. Any urgent request for money or sensitive action that appears to come from an executive should be verified through a separate, trusted channel, never acted on from the request alone. Make this an explicit rule, not an assumption.

AI-Generated Phishing and Fraud

The old advice to spot phishing by its poor grammar and awkward phrasing no longer holds. AI produces grammatically perfect, contextually aware messages tailored to a specific recipient, referencing their real role, colleagues, and recent activity, at massive scale. AI can also make synthetic identity fraud more effective, generating convincing documents, images, and communications around fabricated or blended identities that are used to open accounts and build credit before a coordinated loss.

What This Changes for Defense

The defenses that matter most are not exotic. They are strong identity verification that does not rely on voice or familiarity alone, out-of-band confirmation for consequential requests, staff trained to expect convincing fakes, and a culture where verifying an unusual request is encouraged rather than treated as an insult. The technology of the attack is new; the discipline that defeats it is not.

The uncomfortable truth: the same AI that makes your business more efficient makes your attackers more effective. Adopting AI and defending against AI-enabled attacks are two parts of the same responsibility.

8. Governing AI in Practice

Governance is not a document that sits on a shelf. It is the set of decisions and guardrails that let your organization use AI without losing control of it. Good AI governance is mostly common sense made explicit, assigned to someone, and actually followed.

The Foundation: An AI Policy People Actually Read

Any organization with meaningful AI use or exposure needs a policy that answers a few plain questions in language employees understand. A policy no one reads protects no one, so brevity and clarity matter more than legal completeness.

- Which AI tools are approved, and for which purposes and roles.
- What data may be entered only into specifically approved AI tools, and what data or uses are prohibited entirely.
- When AI output must be reviewed by a person before it is used, especially anything customer-facing or decision-related.
- Who to ask when an employee is unsure, and how to report a mistake without fear of blame.

A One-Page AI Policy Starter

The following is a plain-language skeleton you can adapt into a working policy. It is deliberately short. Fill the brackets with your own specifics, keep it to a single page, and make sure every employee actually reads it. This is a starting point, not legal advice; have counsel review the final version if your industry warrants it.

AI Use Policy [Organization Name]

Purpose. This policy explains how we use AI tools responsibly, so we get their benefit without putting our data, our members or customers, or our organization at risk.

Approved tools. You may use the following AI tools for work: [list]. Using other AI tools for work requires approval from [owner].

Match sensitive data to approved tools. Sensitive data may be entered into an AI system only when that specific tool and that specific use have been approved for that type of information. Never enter customer or member information, account numbers, financial records, credentials, personnel data, confidential or regulated information into an unapproved AI tool. When in doubt, do not enter it, and ask.

Always keep a human in charge. AI may help you draft, summarize, and analyze. A person must review AI output before it is used, and a named person owns any decision AI informs. “The AI produced it” is never an acceptable explanation for a mistake.

Customer-facing and decisions. AI output that reaches a customer or member, or that influences a decision about them, must be reviewed by a qualified person before use. [Add any use that is prohibited outright.]

When unsure, ask. If you are not sure whether an AI use is allowed, ask [owner] before proceeding. If you realize sensitive data went into an AI tool by mistake, report it to [owner] immediately. You will not be punished for reporting a mistake promptly.

Owner and review. This policy is owned by [name/role] and reviewed [quarterly / annually].

Not All AI Use Deserves the Same Governance

A common governance mistake is treating every AI use the same way, which either smothers harmless uses in review or waves through consequential ones. Match the level of oversight to the stakes of the decision.

Low consequence	Brainstorming, internal drafting, summarizing, searching your own documents. Normal employee review is enough. Do not bury these in process.
Moderate consequence	Customer-facing communications and internal analysis that drives action. Requires review and approval by a qualified person before use.

High consequence	Lending, hiring, pricing, legal, safety, security actions, and significant financial decisions. Requires formal approval, validation, monitoring, documentation, and an accountable owner.
-------------------------	--

The Guardrails That Matter Most

Keep humans accountable for consequential decisions. AI can inform a decision that matters, but a named person should own the outcome. The rule is not “no AI in decisions.” It is “no AI decision without an accountable human.”

Make human review meaningful. A human in the loop is not much of a control if that person lacks the expertise, time, information, or authority to challenge the AI. For consequential uses, the reviewer must be capable of recognizing a bad result and empowered to reject or override it. A rubber stamp is not oversight.

Protect your data at the entry point. The most effective control is preventing sensitive data from entering ungoverned tools in the first place, through clear rules, approved alternatives, and, where warranted, technical controls.

Keep an inventory of where AI is used. You cannot govern what you cannot see. Maintain a simple, current list of where AI operates in your business, including AI embedded in vendor tools. At minimum, track the tool or system, the business use, the owner, the data involved, the vendor, and the risk tier. This is one of the most valuable governance artifacts.

Require explainability for anything that matters. If AI influences a consequential outcome, you must be able to explain how, to the level a regulator or a customer would expect. If you cannot, that use is not ready.

Put AI vendors through third-party risk review. Treat an AI vendor like any other party with access to your data. Diligence, contract terms, and ongoing monitoring apply.

Treat approval as a lifecycle, not an event. An AI use that was acceptable when approved may change as the vendor updates the model, the data changes, new capabilities are added, or your organization expands how it is used. Material changes should trigger review rather than inheriting approval automatically.

Know your stop conditions. For consequential AI, define in advance what level of error, drift, security failure, unexplained behavior, or harmful outcome would cause you to suspend or retire the system. Monitoring only matters if it triggers a decision.

Agents Raise the Stakes

Most AI tools produce output: a draft, a summary, an answer. A newer category, AI agents, does not stop at output. AI agents can use tools and take multi-step actions on your behalf: sending messages, modifying records, calling other systems, approving workflows. That shift changes the governance question fundamentally. It is no longer only what the model can say, but what it is authorized to do, because an incorrect output can now become an incorrect action before anyone reviews it.

The controls are recognizable, just applied with more care. Limit what an agent is permitted to do to the minimum its job requires. Define the points where a human must approve before the agent acts, especially for anything consequential. Log the actions an agent takes so they can be reviewed and traced. And keep consequential actions under meaningful human control rather than letting the agent close the loop on its own. Agents can also be manipulated by hostile instructions hidden in the emails, documents, or webpages they process, which is another reason their permissions should stay narrow. NIST has begun treating agent security, identity, and authorization as a distinct area of concern, a signal that this is where a great deal of AI risk is heading. [9]

The Agent Question

Before you let an AI agent act rather than just advise, ask: what is it authorized to do, what must a human approve first, and can we see, stop, and where possible reverse what it did? If you cannot answer all three, the agent is not ready for anything that matters.

Who Owns AI Governance

AI governance fails when it belongs to everyone and therefore no one. Someone at the leadership level must own it, with clear escalation to the board or executive team for consequential decisions. This does not require a new department or a Chief AI Officer in most organizations. It requires a named owner, a simple policy, an inventory, and a regular review. Ownership should be clear, but governance should be cross-functional: business, technology, security, legal, compliance, and privacy should have defined roles where they are relevant. The scale of the apparatus should match the scale of your AI use, but the ownership must be explicit regardless.

Governance Right-Sizing

A small organization does not need an elaborate AI governance program. It needs a short policy, a named owner, an inventory of where AI is used, and the discipline to keep a human accountable for anything that matters. Start there and expand only as your AI use grows.

9. AI and the Regulators

If you operate in a regulated industry, and financial institutions are among the most regulated, AI does not exist in a separate legal space. It is governed by the rules that already apply to you. Understanding that is the difference between confident adoption and an unpleasant surprise at your next examination.

The Core Principle: No New Rulebook, Existing Lenses

For credit unions, the National Credit Union Administration has not issued AI-specific rules or regulations. Instead, it currently supervises AI within its existing framework, applying technology-neutral requirements and focusing on the risks introduced by the use. [1] In practice, AI is governed through the lenses the regulator already examines: information security under Part 748, third-party and vendor risk, fair lending, prohibitions on unfair or deceptive practices, and anti-money-laundering obligations, with recognized frameworks such as the NIST AI Risk Management Framework [6] available to help organize risk management.

One point worth stating plainly, because vendors sometimes get it wrong: NCUA does not currently impose the banking agencies' comprehensive model-risk framework on credit unions. The Government Accountability Office has criticized NCUA's existing model-risk guidance as limited in scope and detail, though NCUA examiners and credit unions may reference banking-agency guidance as a best practice where appropriate. [2] The absence of a dedicated AI model-risk rule is not the absence of expectations. The existing frameworks apply with real force.

What a Defensible AI Program Should Be Able to Show

NCUA supervises AI within its existing risk-focused framework rather than through a separate AI examination. [1] [3] Its current AI resource guidance emphasizes third-party due diligence, internal controls, ongoing monitoring, and understanding the risks introduced by AI-enabled

vendors, linking back to long-standing third-party guidance. [4] None of that requires a specific document in a specific form. But a credit union that can show the following is in a far stronger position, whatever an examiner asks:

- Do we have a documented AI policy or governance standard, approved at the level appropriate to our governance structure?
- What AI tools and systems are in use, including AI embedded in third-party platforms?
- How do we evaluate AI vendors within our third-party risk management program?
- If AI influences lending decisions, how have we validated that outcomes are lawful, consistent, and free from prohibited discrimination?
- Who owns AI governance in management, and how is consequential AI risk reported to the board?
- What is our incident response plan for an AI-related failure or data exposure?
- How do we monitor AI systems for accuracy, bias, and drift over time?

If you are not a credit union, the specifics differ but the shape does not. Whatever regulator governs you, the pattern is the same: AI is examined through existing obligations around fairness, data protection, vendor risk, and accountability. The institutions that do well are the ones that can show a policy, an inventory, and a human owner.

The Fair-Lending and Fairness Trap

One of the highest-risk regulatory areas is any AI that influences who gets credit, at what price, or on what terms, and by extension hiring, pricing, and similar consequential decisions. If an AI system produces outcomes that disadvantage people on a prohibited basis, the fact that a model made the decision is not a defense. The outcome was produced systematically and at scale, which is precisely why it must be validated. Any AI influencing these decisions should be validated for lawful, consistent, and explainable outcomes, with documented controls against discriminatory treatment and a human accountable for the results. [5] And be careful whose validation you are relying on: if the only evidence a model is fair comes from the vendor that sells it, you do not really have independent validation. This is the use case where “a decision you cannot explain is a decision you cannot defend” carries the most weight.

10. Questions Leaders Should Ask

These questions create a practical governance cadence around AI. They require no technical expertise to ask, and the answers reveal whether your organization is adopting AI deliberately or accumulating it by accident.

1. What AI tools are our employees actually using right now, sanctioned or not, and what data has gone into them?
2. Where does AI already operate in our business through tools and vendors we already have?
3. Do we have a written AI policy that employees have read, and a named owner for AI governance?
4. For any AI that influences a decision that matters, who is the accountable human, and can we explain how it works?
5. What sensitive data are we confident is not being entered into ungoverned AI tools?
6. Have our AI vendors been through the same third-party risk review as our other critical vendors?
7. If AI influences lending, hiring, or pricing, have we validated that its outcomes are lawful, consistent, and free from prohibited discrimination?
8. Are our people trained to recognize AI-enabled fraud such as deepfakes, and do we have a plan if an AI tool fails, leaks data, or produces a harmful outcome?
9. If we have or are buying an AI agent, what is it authorized to do, what must a human approve first, and can we see, stop, and reverse what it did? And are we buying it because the work genuinely requires action, or because a vendor renamed a chatbot?
10. Are we adopting AI where it creates real value, or because we feel pressure to be seen using it?

11. Getting Started Without Getting Burned

The path to real AI value is not a dramatic transformation. It is a series of deliberate, low-risk steps that build capability and confidence while keeping control. The following sequence works for most organizations.

First, see what you already have. Before adopting anything new, find out where AI is already in use, in employee habits and in vendor tools. You cannot govern or improve what you cannot see. This inventory is step one and it usually surprises people.

Second, put a simple policy and an owner in place. A short, readable AI policy and a named person responsible for it. This does not need to be elaborate. It needs to exist and be followed.

Third, protect your data. Make crystal clear what information may not enter unapproved AI tools, which approved environments may handle sensitive information, and under what conditions. Give people approved alternatives so the rule is realistic. This closes one of the most immediate exposures.

Fourth, start with low-stakes internal wins. Apply AI first to internal tasks your people already understand well enough to catch mistakes: drafting, summarizing, searching your own documents. Build fluency where errors are cheap.

Fifth, measure the real return. Track what actually changed: time saved, rework required, error rates, adoption, quality, and how much human review the tool still needs. This is the same net-of-review test from earlier, applied to your own results rather than a vendor's claims.

Sixth, govern before you go consequential. Only move AI into customer-facing or decision-making roles once you have the policy, the human accountability, the monitoring, and the explainability to do it responsibly.

This measured path captures meaningful value while reducing avoidable risk. A common failure pattern is reversing the order: deploying AI into consequential or customer-facing decisions first, discovering the risks second, and building governance third, after something went wrong.

Getting the Order Right

You do not have to choose between adopting AI and staying safe. You have to do them in the right order: see what you have, govern it simply, protect your data, win where the stakes are low, and earn your way up to the consequential uses. Deliberate beats fast, and deliberate is also faster in the end, because you are not cleaning up avoidable damage.

When AI Goes Wrong

Have a short plan for the day something breaks. Stop or isolate the affected use. Preserve the prompts, outputs, logs, and relevant records. Revoke access or credentials if an agent or integration is involved. Determine whether sensitive data was exposed or a consequential decision was affected. Then route the event through the same legal, security, privacy, compliance, and notification processes you would use for any other material incident. AI incidents are not a separate world; they are incidents.

Closing: A Tool, Used Well

AI is not the end of business as we know it, and it is not a passing fad. It is a genuinely powerful tool that, like every powerful tool before it, rewards the people who understand what it is good for and respect what it is not.

For business leaders, the task is not to become a technologist. It is to stay clear-eyed while others are swept up in either hype or fear. Adopt AI where it creates real value. Govern the risk it introduces. Keep a human accountable for anything that matters. Protect your data as though it cannot be recovered once it leaves, because often it cannot. And judge every AI claim by what it actually does, how you will know when it is wrong, and who is responsible when it is.

The organizations that will do best with AI are not the ones that adopted it fastest or feared it least. They are the ones that adopted it deliberately, governed it honestly, and never confused a confident answer with a correct one.

AI is a remarkable assistant and a poor authority. Build everything you do with it on that one sentence, and you will get the value while avoiding most of the harm.

Sources and References

The regulatory statements in this guide are drawn from the following primary sources. AI guidance is evolving quickly, so leaders should confirm current expectations with their own counsel and their regulator before relying on any specific point for a compliance decision.

[1] NCUA, Artificial Intelligence (AI) Regulatory and Compliance Resource Center (last modified April 2026). States that NCUA has not issued AI-specific rules or regulations, that existing technology-neutral regulations apply to AI use, and that NCUA supervises AI within its existing framework with a focus on risk management rather than the tool itself.

[2] U.S. Government Accountability Office, “Artificial Intelligence: Use and Oversight in Financial Services” (GAO-25-107197, 2025). Finds NCUA’s model risk management guidance limited in scope and detail, notes NCUA lacks authority to examine technology service providers, and recommends both be strengthened.

[3] NCUA, 2026 Supervisory Priorities (Letter to Credit Unions 26-CU-01, January 2026). Establishes the agency’s current risk-focused supervisory priorities and examination-program focus.

[4] NCUA, Letter to Credit Unions 07-CU-13, Evaluating Third Party Relationships (2007), and Letter to Credit Unions 01-CU-20, Due Diligence Over Third Party Service Providers (2001). Both remain active and are specifically referenced by NCUA’s current Artificial Intelligence Resource Center for evaluating AI vendors.

[5] NCUA, Letter 25-CU-04, Removal of Disparate Impact (2025). Instructs examiners to discontinue requesting or reviewing internal disparate-impact analyses, while continuing fair-lending supervision focused on applicable law and evidence of disparate treatment.

[6] National Institute of Standards and Technology, AI Risk Management Framework (AI RMF 1.0) and NIST AI 600-1, the Generative Artificial Intelligence Profile. Voluntary resources for organizing and managing AI and generative-AI risk.

[7] Financial Crimes Enforcement Network (FinCEN), Alert on Fraud Schemes Involving Deepfake Media Targeting Financial Institutions (2024). Warns of increased suspicious activity involving GenAI-generated or altered identity materials used to evade verification.

[8] U.S. Copyright Office, Copyright and Artificial Intelligence, Part 2: Copyrightability (2025). Maintains that human authorship remains central to copyright protection for AI-assisted and AI-generated material.

[9] NIST, Center for AI Standards and Innovation, AI Agent Standards Initiative (announced February 2026), and related NCCoE work on AI agent identity and authorization. Treats the security, identity, and authorization of autonomous AI agents as a distinct area of concern.

About the Author

J.W. Walker has spent his career in information security across commercial and government environments. He has sat in the tester's chair, built a company from the ground up through exit, and serves as CTO of an information security firm. Today he advises executives and boards on technology strategy, security, and the practical adoption and governance of emerging tools, including AI. He writes plain-language guides that help leaders make confident decisions about technology and risk without needing to become technical experts.

If you found this guide useful, visit [**jeremywalker.us/resources**](https://jeremywalker.us/resources) for more guides on cybersecurity, AI, and technology leadership.

The Business Leaders Series

Plain-language guides for the decision-maker, not the technician. Each one takes a subject that is usually buried in jargon and makes it something a leader can actually use. If this guide was useful, the others are written the same way.

Penetration Testing for Business Leaders

What to test, what to ask, and how to turn findings into action.

Cybersecurity for Business Leaders: CU Edition

Cyber risk, oversight, and member protection, written for the board.

Cybersecurity for Business Leaders: SMB Edition

Cyber risk, protection, and practical steps for owners. Written with Ideal Technologys.

AI for Business Leaders *(you are reading this one)*

Getting real value from AI without getting burned by it.

The Virtual CISO for Business Leaders

What you are really buying when you hire fractional security leadership.

Personal Cybersecurity for Business Leaders

Protecting your identity, money, devices, family, and digital life. Built to share with your team and family.

Find the complete series at **jeremywalker.us/resources**

This series is refreshed periodically. See the sources page of each guide for its last-reviewed date.

Last reviewed: August 2026. Regulatory positions and standards evolve; confirm specific requirements with counsel or your regulator before relying on them.