



The Human and AI Company

How to reorganize around the customer, decide who owns AI, and move your people from doing the work to directing it.

Your customers can't feel your AI yet.

Everyone wants to talk about generative AI. Boards ask about it, vendors pitch it every week, and most companies can now point to a chatbot, a stack of licenses, and a handful of pilots. What very few can point to is a customer who got a faster answer, a cleaner order, or a problem fixed the first time because of any of it.

That gap has an explanation. Most companies run AI as a writing assistant that drafts the email, summarizes the meeting, and produces a first version of the proposal. The individual task gets faster, but the customer's request still travels through the same departments, waits in the same approval queues, and stalls at the same handoffs, because nothing about the way the work is organized has changed.

The real power of AI shows up when it carries a piece of work from the first request to the finished result, with people setting the rules, handling the exceptions, and answering for the outcome. A company that gets there can respond in hours where it used to take days, and it can grow without adding people every time volume rises.

88% / 23%

USE AI / SCALE AI AGENTS

In McKinsey's 2025 global survey, 88% of companies said they use AI in at least one business function, but only 23% were scaling AI agents anywhere in the business. The companies getting the most value were **nearly three times as likely to have fundamentally redesigned their workflows.**¹

So the obstacle is rarely the technology, which every one of your competitors can buy from the same vendors you can. The obstacle is the way we do business today: organizations built around departments, decisions that sit with whoever has always made them, and nobody who owns the customer's result from one end to the other. That has to change before AI can do much more than type faster.

This paper lays out how to change it. The eight steps that follow begin with a single customer job and end with an organization in which AI handles the routine flow, your people direct it, and your customers can tell the difference.

IF YOU READ NOTHING ELSE

Start with what the customer receives, and pick a job you can watch from request to finish.

Build the team around the flow: agents run it, one person owns the result, specialists take the exceptions.

Let AI earn its independence through measured results, one type of decision at a time.

WHO THIS IS FOR

Chief executives and operating leaders of mid-size companies who have tried AI, have little to show customers for it, and want a practical way to reorganize around it.

From people doing the work to people directing it

It helps to see the destination before choosing the first step, because the pilot you run this quarter should be the first rung of a ladder and not a one-off experiment.

STAGE 1

AI assists

AI helps a person do a task, and the person still does the work. Individual productivity improves, the customer notices very little, and this is where most companies sit today.

STAGE 2

Human in the loop

AI does the work and a person approves each result before it reaches the customer. The work gets faster, although the approver soon becomes the new queue.

STAGE 3

Human on the loop

A small team directs a set of agents that run the flow. One person owns the result, specialists take the exceptions and keep the agents sharp, and anyone on the team can stop it.

Stage three is where the economics change, since routine work no longer waits for a person to become available and your most experienced people spend their time on the cases that need judgment. It is also where the risk concentrates, which is why no AI should start there. Every agent begins with a person approving its work, and it earns more freedom only by showing, in measured results, that it performs at least as well as the people it supports.

None of this is a technology project. Your competitors can buy the same models and the same software. What they cannot buy is an organization that knows who owns each customer result, who owns each AI agent, which decisions AI may make alone, and how that permission is earned and withdrawn. Building that organization is the subject of this paper.

Eight steps

01

Find the job

Recurring work with a clear finish and a real delay.

02

Give AI a role

A written assignment with inputs, outputs, and limits.

03

Organize around it

A small team directing named agents, with one owner for the result.

04

Decide who owns AI

Central control, distributed building, one owner per agent.

05

Bring your people

Workload, status, and skills settled before launch.

06

Set controls

Review that works and the budget to sustain it.

07

Earn autonomy

Decisions sorted by risk and promoted on evidence.

08

Prove, then scale

A staged rollout and a fair comparison before you expand.

Find the job worth changing

Start with a piece of work that recurs and has a clear finish, such as a proposal ready to send, a service request resolved, or an account reconciled, and give it an owner who has the authority to change how it runs.

Define the finish from the customer's side

Writing a reply is a task, while getting the customer an accurate answer is the job. If your pilot ends at an internal approval, name who carries the work through the next handoff, because a faster internal step that still leaves the customer waiting has improved nothing they can see.

- 01 Walk real cases.** Sit with the people who do the work and go through several ordinary cases and at least one that went wrong.
- 02 Find the wait.** Note every point where work stalls: missing information, a correction, an approval queue, or a handoff.
- 03 Separate working time from waiting time,** and mark which numbers were measured and which estimated.
- 04 Check that you can test it.** You will need real examples, someone who knows the job well, and enough volume to compare before and after.

Aim at the constraint

90 min / 3 days
WRITING TIME / WAITING TIME

Suppose a proposal takes 90 minutes to write and then waits three days for a discount approval. You could cut the writing time in half and **the customer would still wait three days**, so the approval is the thing to fix.

WHAT YOU FIND	WHERE TO LOOK FIRST
The team keeps going back to the customer for details	Intake, and in particular who confirms requirements and when
Drafts are built from scattered information	Approved sources and the people who maintain them
Routine decisions wait on one person	Approval limits, backup approvers, and ownership of the queue
Finished work comes back	Quality checks and responsibility across handoffs


Speeding up one department often just builds a queue in the next, so measure the whole job before deciding where AI belongs. Decide in advance how you will use the time you free, whether to clear a backlog, take on demand you now turn away, or cut overtime, because freed time with no plan never shows up as a saving.

Give AI a defined role

Lay out the full sequence of the job before assigning any of it. Drop the steps that add nothing, simplify the handoffs, and put repeatable calculations in rules or a spreadsheet. AI earns its place where information has to be sorted, a draft prepared, or several steps coordinated.

Can AI run the job end to end?

It can, provided your tools, information, and system connections support it, and provided you keep running the sequence separate from holding authority over it. Approving a price, committing a date, and sending the proposal each need explicit permission and a control behind them, so begin with AI preparing the work and a person receiving it, and treat anything beyond that as something to be earned in Step 07.



Name the agents and give each one a job description

"AI" is not a role. An agent is a named piece of software with one job, one owner, defined inputs, and written limits. The proposal flow needs three.

AGENT	ITS JOB	IT MAY USE	IT MAY NOT
Intake agent	Collects the request, checks the required fields, asks the customer for what is missing	The intake form and the CRM	Promise anything
Drafting agent	Writes the scope from confirmed requirements and the approved service list, showing sources and open questions	Current terms, pricing rules, customer facts	Invent dates, approve exceptions, send
Status agent	Tracks each proposal through the checks, routes it to the approver, flags anything stalled	The queue and the approval rules	Approve or skip a check

Judge each handoff by what the person receiving it has to do: if the specialist must rebuild the request to check the draft, the work has moved into review, not been saved.

Test AI against the simpler fix

Compare three options: today's process, a better version without AI (cleaner intake, a template, a clearer approval limit), and the AI design. When the simpler fix wins, take it.

106

EXPERIMENTS REVIEWED

A 2024 review in *Nature Human Behaviour* found human and AI combinations performed, on average, **worse than the better of the two alone**. Pairing helped in creating content and hurt in making decisions.² Where you place AI matters more than whether you use it.

Organize around the job


Departments exist to group expertise in sales, operations, finance, and technical support, but the customer's job cuts across all of them. AI can smooth the handoffs between departments, yet somebody still has to hold authority over the whole result.

The way to get there is to design backwards. Start from the finished job, work out the skills, decisions, and tools it requires, and let the departments supply them while the design follows the work.

Does this break up the silos?

It changes what the work is and where decisions sit. The agents do the routine steps that used to pass from desk to desk. A proposal owner directs them, resolves what they cannot, and follows each job across departments, while department managers commit the people and the hours that this requires.

Where the volume justifies it, you can go further and move people into a team that owns the whole job, keeping departments wherever their expertise, career development, or controls still earn their place.



Managers keep managing their people unless you change the reporting lines. The agents report to nobody; they have an owner.

Keep reporting and workflow separate

Reporting

Who manages workload, performance, development, and employment decisions?

Work ownership

Who follows the customer job, resolves handoffs, and approves process changes?

Decision rights

Who can accept scope, approve a price exception, commit capacity, or stop a release?


Priority conflicts

Who decides when two workflows need the same expert at the same time?

Two cautions apply. An AI agent that routes tasks is nobody's manager, and the title of workflow owner gives a person no control over another department's capacity, so both arrangements need to be agreed with the reporting managers.

The team that runs an AI-enabled flow

This is the working unit of the company you are building. A handful of people direct a set of agents that run the flow. The people are not doing the routine work any more; they are setting the rules, taking the exceptions, and answering for the result.



Workflow owner

A generalist who owns the flow from end to end. Watches how the agents perform, triages what they escalate, adjusts the limits, and can stop the line without asking. Your best coordinator or planner grows into this.

Specialists

The experts who take the cases agents cannot handle, correct the agents when they are wrong, and rebuild the process when one keeps failing. They are the escalation bench and the insurance policy.

Technical owner

Keeps the agents connected, monitored, and versioned, and recovers them when they fail. Owns agent health and never the business result.

Department managers

Still manage the people: workload, development, performance, and employment decisions. They commit hours to the flow and agree the limits with the owner.

The specialist bench is the part leaders are tempted to thin first, because once agents run the flow the experts look underused. Judging an agent's output takes the same depth as doing the work, so when the bench goes, the ability to catch errors goes with it, and the first sign is usually a customer. Protect it and keep it practicing.

Two shapes of team, three places to put it

The mix of people and agents follows the nature of the work, and the team can sit inside your departments or alongside them. Choose both deliberately.

The shape follows the work

A stream of volume

Orders, service requests, scheduling, routine proposals, invoicing. Agents handle the volume, one owner supervises it, and a small specialist bench is on call for whatever the stream throws off. People are a minority of the team's capacity, and the owner's job is mostly exceptions and limits.

A set of relationships

Key accounts, major bids, complex negotiations, strategic suppliers. Specialists lead each case, agents prepare the analysis and the paperwork and close the small, routine instances on their own. People remain the center of the team, and the agents are their staff.

Most companies have both. A professional services firm runs its delivery-to-cash work as a stream and its biggest client relationships as a set, and the two teams can share the same specialists.

Where the team sits

Inside departments

People stay in sales, operations, and finance; the workflow owner directs the flow across them and the agents belong to the flow. Start here, because it is cheapest and easiest to reverse. Watch for an owner who can chase work but cannot get capacity or a decision.

A dedicated team

Move the people onto a team that owns the whole job when demand is steady and the handoffs keep failing even with an owner in place. Agree who manages development and how expert standards stay consistent.

Shared specialists

When several flows lean on the same expert or the same tools, run that expertise as a service with an owner, a queue, allocated hours, and standards. Watch for agent output burying the shared reviewer.

Size changes the form these take. In a smaller company the existing leads double up and bring in outside help, while a mid-size firm needs explicit agreements between department heads and workflow owners. Let the work, and not a headcount threshold, tell you when to restructure.

THE TEST OF A REAL REORGANIZATION

Redrawing the org chart achieves nothing by itself. The change is real when three things are true: named agents are doing the routine steps, one person owns the result and can stop the flow, and the specialists spend their time on exceptions and corrections instead of the routine work.

Decide who owns AI

Most companies drift into one of two arrangements, and neither holds up. Either a central team owns every AI project and becomes a bottleneck, or each department builds its own and nobody can say what is running or who answers for it.

What works is to keep the rules in one place and the ownership in the business. One accountable leader sets the standards for tools, security, and testing, while the people who run the work decide where AI is used and answer for results. In an IBM study of more than 600 chief AI officers, organizations that ran a centralized or hub-and-spoke model under an accountable AI leader reported up to 36% higher returns on their AI initiatives.³

Five responsibilities across the team

RESPONSIBILITY	WHAT THE PERSON OWNS
Workflow	The customer result, queues, handoffs, exceptions, and process changes
Domain expertise	Standards, approved knowledge, hard judgment calls, and expert feedback
Agent lifecycle	Each agent's purpose, instructions, sources, tests, updates, measured value, and retirement
Technical operation	Connections, permissions, monitoring, versions, recovery, and exit from a provider
People development	Practice, feedback, progression, performance expectations, and workload

74% / 21%

EXPECT AGENTS / CAN GOVERN THEM

In Deloitte's 2026 survey of 3,235 leaders, 74% expected their organizations to be using AI agents at least moderately by 2027, while **only 21% said they had a mature model for governing them.**⁴

Three rules that keep ownership real

- 01 Every AI agent has one named owner**, who is a person and not a team. If nobody will put their name to an agent, it does not go live.
- 02 Accountability stays with the business.** The manager who sets what an agent may and may not do answers for what it does inside those limits, in the same way they answer for the work of their team. Technical staff keep the system running but do not own the commercial result.
- 03 People need authority they can use on an ordinary day.** Spell out what a specialist can accept alone, set approval limits for exceptions, let a qualified reviewer stop a step, and name one leader who settles conflicts between workflow owners and department managers.

Bring your people with you

Talk to the people who do the work before anything is rolled out, and ask them what gets easier, what gets harder, and where the plan creates extra work. When someone tells you the review queue is impossible, they have found a flaw in the design, and the design is what should change.

Settle five things before launch

Purpose	Explain the customer result you are after, what is decided, what is still open, and how their feedback will change the pilot.
Workload	Set aside time for review and learning, and say plainly what leaves the old task list.
Incentives	Remove the mixed messages, since you cannot reward billable hours while asking people to reduce them.
Latitude	Give people real choices inside the limits, practice that counts, and access to expert feedback.
Fairness	Explain what is monitored and how people are evaluated, and give staff a way to report errors or disagree that protects anyone who uses it in good faith.

Promise only what you intend to honor on jobs and pay, and tell people how freed time will be used; if that has not been decided, say so. The research on motivation and psychological safety is consistent: people do better work and raise problems sooner when they have a say, the skills to succeed, and no fear of speaking up.^{5, 6}

Keep expertise from eroding

Judging AI's work takes the same depth of knowledge as doing it, and that knowledge fades when people stop using it. Students who practiced with an unrestricted AI tool scored 17% lower than classmates once it was taken away,⁷ and a trial with junior software engineers found a similar gap.⁸

28% → 22%

DETECTION RATE WITHOUT AI

Experienced professionals are not immune. At four endoscopy centers, doctors' detection rates in procedures performed without AI **fell from 28.4% to 22.4% within months of AI becoming routine.**⁹ The loss stays hidden until the day the automation is unavailable.

Practice

An ordinary case



Explain

The choices made



Feedback

From the expert



New case

Without AI

On the proposal team, a junior drafts a low-risk case, compares it with the AI version, goes through the scope mistakes with the estimator, and then checks a changed request with no assistance. Budget for the learner's time and the expert's, and schedule regular practice without AI for experienced staff as well, so that the team can still run the job by hand when it has to.

Set controls for the real job

Write down what AI may do, what information it may use, what gets checked, and what stops the work, and then build those limits into the tool setup and the team's instructions so that they do not depend on anyone's memory.

A control plan for the proposal job

CONTROL	WHAT TO AGREE
Information	Confirmed requirements together with the current service list, terms, and pricing rules, with named owners who retire old versions
Permissions	AI drafts and flags gaps, and it does not send, grant price exceptions, or commit delivery dates without explicit authorization
Release checks	The specialist checks scope, approved rules calculate the price, operations confirms capacity, and authorized people approve and send
Stop and recovery	Missing evidence or an unsupported promise pauses the step, a named person investigates, and a qualified backup runs the manual process

A review is a control only when the reviewer has four things

Expertise They know the work and the errors that matter.	Evidence Current source information sits beside the output, independent of it.	Time Enough hours for the volume, including exceptions and corrections.	Authority They can reject the work, ask for information, or stop the step.
--	--	---	--

Use formulas wherever an answer can be verified and qualified people wherever judgment is needed. Two AI models agreeing with each other proves nothing, and neither does an explanation that happens to read well.

Then test the control itself. Plant known errors under realistic conditions and record what is missed and what is flagged wrongly. Run ordinary work, difficult work, and a failure serious enough to stop the pilot, having named in advance who pauses it and who fixes it, because consequential mistakes have to be caught before release and cannot reliably be undone afterward.

319

KNOWLEDGE WORKERS SURVEYED

The more confidence people had in the AI, **the less critical thinking they reported applying to its output.**¹⁰ Skepticism cannot be left to individual vigilance, so every check should be a named step with an owner, and a reviewer who never overrides anything should be treated as a warning sign.

Fund the review, the support, and the upkeep

Faster preparation builds a bigger queue for whoever checks the work, so review capacity has to be confirmed before AI is asked to produce more.

15 of 60

OUTPUTS REVIEWED PER DAY

Sixty AI outputs a day at eight minutes of review apiece comes to **eight reviewer-hours**. A manager with two hours to spare gets through 15 of them, and the other 45 either wait or go out unchecked. This arithmetic is why approving every item cannot be the end state. Illustrative figures.

Keep a record for every agent

Each agent should have a written record of its purpose, its owner, its sources and permissions, its version, the tests it has passed, who supports it, and the value it has actually delivered. Name who updates its information and who approves significant changes, agree which small changes can take a lighter check, and retire any agent that duplicates work or costs more to maintain than it returns.

Work with the records people already trust

- 01 Start with this job's sources** and reconcile the conflicting service lists and terms. You do not have to clean every document in the company to test one job.
- 02 Name the official record**, including where inputs come from, where the draft goes, and which approved version is final.
- 03 Check what you already own.** A template, an approved export, or a controlled manual transfer may be enough, provided you count the transfer time.
- 04 Confirm access** by checking authorization, the way the provider handles and retains data, and each service involved before any business or customer data goes in.
- 05 Plan for failed connections** by deciding how duplicates and retries get checked, since a request that timed out may still have gone through.

Plan for the day it fails

Set support hours, a backup, a customer escalation route, and a fallback that has actually been tested. If drafting fails, someone needs to know which manual template to pick up; if the specialist is out, the backup needs to have the hours; and if timing slips, someone has to tell the customer.

Budget for upkeep, practice, incidents, and the cost of leaving a provider, and keep instructions and records exportable so that leaving remains possible.

Let AI earn its autonomy

Having a person approve every AI output is the right way to start and the wrong place to stay, because the approver becomes the bottleneck and, under pressure, approval turns into a rubber stamp. The way forward is to sort decisions by what a mistake would cost and let AI earn independence where a wrong answer is cheap, visible, and easy to reverse.

KIND OF DECISION	EXAMPLES	OVERSIGHT	BEFORE AI GETS MORE FREEDOM
Low risk and easily reversed	Status updates, moving data between systems, pulling details from documents, reminders, routine reports	AI acts alone and a person audits a sample	It matches or beats people on real cases, and its errors are cheap to undo
Routine, inside set limits	Standard quotes at list price, routine customer replies, scheduling within agreed rules, invoices that match the contract	A person approves each one at first, and later AI acts while a person can override	A sustained record of accepted work, few overrides with each one explained, no serious incidents, and a tested way to undo
Judgment and consequence	Discounts, unusual scope, refunds above a limit, strategic bids, anything that commits the company in a new way	AI prepares the case and a person decides	Stays with people

Four rules for extending trust

- 01 Every agent starts under approval.** Independence is granted in writing, one type of decision at a time, and withdrawn when an agent steps outside its limits.
- 02 An approval rate is not an accuracy rate.** A reviewer who approves everything may simply have stopped looking, so check a sample of the AI's decisions against what actually happened.
- 03 Read the overrides.** Have reviewers record why they changed an output. A trickle of explained overrides is healthy, none at all is a warning, and a flood means AI was trusted too early.
- 04 Judge the whole flow.** When work passes through several AI steps, small error rates multiply, so measure the finished result and not only each step.

90% × 3 = 73%

THREE GOOD STEPS, ONE WEAKER RESULT

Three steps that are each right 90% of the time give a correct result only about 73% of the time in sequence. **A flow can fail while every part looks acceptable**, so measure the customer's finished result.

Prove it, then scale

Record how the job runs today before changing anything, and choose the cases, the comparison, the quality targets, and the stop conditions before you see any results. The hard cases and the failures stay in the count.

Go live in stages

FIRST

Run it alongside

AI produces its version while people keep doing the job as usual, and you compare the two before any customer sees AI's work.

THEN

Approve each one

AI's work goes out once a person has approved it, and every change the reviewer makes is recorded with a reason.

LATER

Supervise

The safest types of case move to AI acting alone with a person able to override, as described in Step 07.

Running AI alongside the existing process takes time and is worth it, because the alternative is discovering the first serious error in front of a customer.

Measure the whole result

Customer

Time to an accurate answer, rounds of clarification, corrections after sending, and issues that come back.

Labor

All of it by role, including prompting, review, corrections, upkeep, practice, and work pushed onto others.

Queues

Waiting time, review delays, overtime, and pressure on shared specialists.

Capability

Whether people can still explain decisions, handle new cases, and run the fallback.

The headline number is **accepted results per total labor hour**, with setup kept separate from recurring work and every hour counted once. Where you can, assign comparable work at random, and where you cannot, run a phased comparison and write down its limits.

Make the call, then choose the next job

There are four honest outcomes, which are to continue, revise, take the simpler option, or stop, and a pilot showing that intake, approval rights, or review capacity must be fixed first has done its job. Before a successful pilot becomes normal work, confirm who owns it, who supports it, and who maintains its sources.

WHAT KEEPS RECURRING	WHAT IT CALLS FOR
Shared experts are always the queue	More review capacity
Handoffs keep failing	A dedicated team for that job
Instructions conflict	Common sources, with one owner for changes

One job, taken up the ladder

A 300-person B2B services company sends about 300 proposals a month. Each takes roughly 90 minutes to prepare and then waits around three days for approval, and customers have started to mention it.

PHASE 1 Redesign The sales leader becomes workflow owner. Intake is standardized and approval limits are written down.	PHASE 2 Approve first Agents draft every proposal from approved sources. The estimator checks scope and approvers sign, with overrides logged.	PHASE 3 Supervise Standard proposals at list price, about 60% of volume, go out the same day. The estimator audits a weekly sample.	ALWAYS People lead Discounts, unusual scope, and strategic bids stay with people, who get a prepared case.
--	--	---	--

WHAT CHANGES	TODAY	AFTER PHASE 3
Customer's wait for a standard proposal	About three days	Same day (the target to prove)
Preparation labor each month	450 hours (300 at 90 minutes)	About 255 hours (300 at 45 minutes, plus 30 for upkeep)
Approval of standard proposals	A manager signs each one	A weekly sampled audit

3 days → same day

WAIT FOR A STANDARD PROPOSAL

The labor saving is modest: 195 hours a month, about \$11,700 at a \$60 loaded rate. **The larger prize is the customer's wait.** Answering in hours while rivals take days is an edge they cannot buy.

Owner	The sales leader directs the agents and answers for the result
Specialists	The estimator and the operations lead take exceptions and correct the agents
Technical	Sales operations keeps the agents connected and versioned
Managers	Department managers still manage their people

AGENT	OWNER	TRUST LEVEL
Intake agent	Sales leader	Acts alone, sampled audit
Drafting agent	Estimator	Acts alone on standard scope; approved on the rest
Status agent	Sales operations	Acts alone, sampled audit

Still unconfirmed: interfaces, permissions, review minutes, and the fallback. Stop test: a planted draft promising a service the firm does not offer; if the check misses it, the drafting agent goes back to approval.

The one-page operating brief

Everything in this paper lands on one page of named people and visible unknowns. If you can fill it in, you are ready to pilot; wherever you cannot, you have found your next piece of work.

01

Customer job

Trigger, finish, owner, baseline, the delay you are going after, and the facts still missing.

02

Task assignments

What each person, rule, or agent does, with inputs, output, and handoff, and the alternatives you compared.

03

The team

The owner, the specialists, the technical owner, and the named agents, with who accepts scope, commits capacity, approves, sends, and stops.

04

AI ownership

The named owner of each agent, the manager who sets its limits, and its written record.

05

People

Changed jobs, work removed, review time, practice, feedback, and who owns learning.

06

Information and tools

Official records, source owners, versions, permissions, provider data handling, and connections not yet built.

07

Controls and support

Error checks, review capacity, escalation, support hours, and a tested manual fallback.

08

Autonomy plan

The risk level of each decision, what AI must show before acting alone, and what withdraws that trust.

09

Pilot and costs

Comparison, cases, stages, targets, total hours, cash costs, stop conditions, and review date.

DILLINGHAM CAPITAL

Want a second set of eyes on your brief?

Fill it in with someone who knows the job, then send it to us, sharing only what the review needs. We will tell you what is ready to test, what is missing, and where setup help is worth paying for.

Request the companion workbook or a brief review at dillinghamcapital.com

ABOUT THE AUTHOR

Aaron Dillingham leads Dillingham Capital. He builds production AI, machine learning, and analytics systems for large-scale operations. He is a former Army combat medic and wildland firefighter, holds an Ed.D. in organizational leadership from Vanderbilt and an MBA from Texas A&M, and wrote *Stop the Bleed: A Former Combat Medic's Framework for Business Crisis Leadership*.

Sources

WHAT THE EVIDENCE DOES AND DOES NOT SHOW

The research supports the design choices in this paper. It does not prove this operating method, and it cannot predict how AI will perform in your company or what return you will see. The worked example uses invented numbers, and some survey figures are self-reported. Your own pilot is the evidence that counts.

- 1 **McKinsey (2025).** The State of AI. Global survey of 1,993 respondents, fielded June to July 2025. [Report](#)
- 2 **Vaccaro, Almaatouq, and Malone (2024).** When combinations of humans and AI are useful. *Nature Human Behaviour*. A review of 106 experiments published from 2020 to 2023. [Paper](#)
- 3 **IBM Institute for Business Value and Dubai Future Foundation (2025).** Chief AI Officer study of more than 600 CAIOs in 22 countries. Self-reported returns. [Release](#)
- 4 **Deloitte (2026).** State of AI in the Enterprise. Survey of 3,235 business and IT leaders in 24 countries. [Report](#)
- 5 **Deci, Olafsen, and Ryan (2017).** Self-Determination Theory in Work Organizations. A research review on autonomy, competence, and relatedness at work. [Paper](#)
- 6 **Frazier and colleagues (2017).** Psychological Safety. *Personnel Psychology*. A meta-analysis on speaking up and reporting problems at work. [Research page](#)
- 7 **Bastani and colleagues (2025).** Generative AI without guardrails can harm learning. *PNAS*. A field experiment with nearly a thousand high-school math students. [Paper](#)
- 8 **Shen and Tamkin (2026).** How AI Impacts Skill Formation. Anthropic. A randomized trial with 52 mostly junior software engineers, in which the AI-assisted group averaged 50% on a follow-up quiz and the hand-coding group 67%. [Research report](#)
- 9 **Budzyń and colleagues (2025).** Endoscopist deskilling after exposure to artificial intelligence in colonoscopy. *The Lancet Gastroenterology & Hepatology*. An observational study at four centers in Poland. [Record](#)
- 10 **Lee and colleagues (2025).** The Impact of Generative AI on Critical Thinking. *CHI*. A survey of 319 knowledge workers on self-reported checking effort and confidence. [Research page](#)

© 2026 Dillingham Capital LLC. This paper offers general guidance for business leaders and is not legal, financial, or employment advice. It draws only on published research, and the views are the author's own and do not represent any employer.