



Computer Instruments

Governing AI Readiness: Business Value With People at the Center

Why AI pays off when it's governed and built around the people who use it, and how to tell whether you're ready

WHITE PAPER • OCTOBER 2026

Governing AI Readiness: Business Value With People at the Center

Why AI pays off when it's governed and built around the people who use it, and how to tell whether you're ready

By Computer Instruments · Published October 29, 2026 · 9 minute read

Executive Summary

Most organizations use AI now. Far fewer can point to what it has earned them. In McKinsey's 2026 survey, nearly nine in ten respondents said their organizations use AI regularly in at least one business function, and eight in ten said it had improved their own productivity. Only 37% said AI had added to their organization's EBIT, about the same share as the year before.¹

A study of 5,172 customer-support agents found that an AI assistant raised the number of issues resolved per hour by 15% on average. The least experienced agents gained the most: they got faster, and their answers got better. Customers in those chats were more polite and less likely to ask for a manager.² At the other end, Gartner expects more than 40% of agentic AI projects to be canceled by the end of 2027 because of escalating costs, unclear business value, or inadequate risk controls.³

Most of the difference between those two outcomes is settled before anything goes live. Did someone measure the starting point? Did anyone decide who owns the tool and what it's allowed to touch? Were the people who do the work part of the design? This paper sets out a model in three parts:

1. **Assess.** Score readiness across eight dimensions, from data and knowledge to governance and how outcomes will be measured, before any money is committed.
2. **Govern.** Give each use case an owner, guardrails, and a review schedule, following the NIST AI Risk Management Framework and ISO/IEC 42001.
3. **Empower.** Use AI to take low-value work off people's plates and bring every employee closer to the performance of your best, with people keeping the judgment calls.

Done in that order, every AI dollar is tied to a measured outcome you can explain to regulators and customers, and employees treat the tools as something that helps them.

Why Governance Comes First

AI is already inside most organizations, whether or not anyone approved it. In Microsoft and LinkedIn's 2024 survey of 31,000 people in 31 countries, 75% of knowledge workers said they use AI at work. Of those users, 78% bring their own tools, and only 39% have had AI training from their employer.⁴ When leadership hasn't set the rules, employees set their own. Customer records, pricing, and internal policy end up in whichever tool someone signed up for.

The cost shows up in three places.

Spending without a baseline. In McKinsey's 2025 survey, fewer than one in five respondents said their organizations track KPIs for their generative AI tools.⁵ A project with no measured starting point can't show what it changed, so it's the first thing cut when budgets tighten.

Unreviewed output. In the same survey, 27% of organizations using generative AI said employees review everything it produces before it's used.⁵ The rest review less than all of it.

Rules with penalties attached. Texas's Responsible Artificial Intelligence Governance Act took effect on January 1, 2026. Government agencies must tell people when they're interacting with an AI system, health care providers must disclose their use of AI in certain settings, and certain uses are banned outright. The Attorney General enforces it, with civil penalties of \$10,000 to \$200,000 per violation and up to \$40,000 a day for violations that continue.⁶ California followed in September 2026 with AB 1609. Starting January 1, 2027, it bars large businesses from presenting a customer service chatbot as a human and requires them to let customers reach a person during business hours.⁷

Governance is also where the value comes from. Of the 25 attributes McKinsey tested in 2025, redesigning workflows had the biggest effect on whether an organization saw EBIT impact from generative AI. Only 21% of organizations using it had fundamentally redesigned even some of their workflows.⁵ In the 2026 survey, the 6% of respondents McKinsey counts as high performers stood out for the same habit, along with human-in-the-loop design, impact measurement, and risk management.¹ Redesigning a workflow takes someone with the authority to decide what the work should look like, who owns it, and how success will be measured.

Where the Value Shows Up

AI pays off in two ways. It handles work that never needed a person, and it helps the person who handles everything else work faster and get it right. Customer service shows both: routine calls can be answered without an agent, and in the study above, agents with an assistant beside them resolved more issues per hour.²

The same split holds in other departments.

Function	Work AI can take	Work AI can help with	What to measure
Customer service	Routine status, payment, and policy calls	Agent answers, call summaries, next steps	Calls resolved without an agent, handle time, first-contact resolution, satisfaction
Sales	Research, call notes, CRM updates	Account briefings, proposal drafts, follow-up	Selling time, pipeline per rep, win rate
Human resources	Benefits, policy, and leave questions	Onboarding, case triage	Ticket volume, time to answer, new-hire ramp time
Finance	Invoice status, collections reminders	Reconciliation, exception review	Days sales outstanding, cost per transaction
IT and operations	Password resets, routine tickets	Outage triage, keeping the knowledge base current	Mean time to resolve, tickets resolved through self-service

The last column is where most organizations fall short: fewer than one in five track KPIs for their generative AI tools.⁵ The ones that get results measure a baseline first, redesign the workflow around the tool, and track the outcome afterward. A readiness assessment gives you that baseline.

AI That Makes People Better at Their Jobs

The most useful finding in the customer-support study is who benefited. The assistant helped newer agents most, because it put the know-how of the best agents in front of everyone.² The goal is AI that brings every employee closer to the performance of your best, rather than AI that replaces them or watches them.

Service leaders are already planning this way. In a Gartner survey of 321 customer service and support leaders, 85% said they're adding new responsibilities to frontline roles as AI cuts contact volume, while only 31% had carried out or planned AI-driven layoffs through early 2027.⁸ Customers want a person within reach, too: 87% told Gartner it's essential for a company using generative AI in customer service to offer a way to reach a human.⁹

Six principles put this into practice.

AI takes the task, and people keep the judgment. Routine work is automated. Exceptions, complaints, and decisions that affect a customer's money, health, or rights stay with a person.

The people who do the work help design it. They point out what slows them down and test the tools before anyone else has to use them.

Training comes before rollout. In 2024, only 39% of AI users had been trained by their employer.⁴ Everyone whose role touches AI learns what it does, where it gets things wrong, and when to override it.

Employees know what's being measured. They know what AI records about their work and how it's used. Coaching data is used for coaching and nothing else.

Freed-up time has a plan. Hours AI gives back go to complex cases, retention calls, outreach, and skills development, as agreed before rollout.

Anyone can raise a concern. An employee who spots a wrong answer, an unfair outcome, or a risky use can flag it, and the flag reaches someone who can act on it.

These principles decide whether employees use the tools you approved or keep bringing their own, and whether the productivity gain ever shows up in your numbers.


The AI Readiness Assessment

The assessment produces a readiness score for each dimension, along with the evidence behind it. It starts with an inventory of every system that touches a customer or employee interaction: the phone system and IVR, call routing, workforce management, quality management, CRM, call recording, transcription, digital channels, and any outsourcer platforms. Each dimension below is then scored on a four-stage scale.

Dimension	What gets checked	Why it matters
Data	Whether interaction and customer data is complete, accurate, and accessible, and who owns each source	AI can only work with what it can read
Knowledge	Whether policies, FAQs, and procedures are current, consistent, and in a form AI can use	Out-of-date knowledge produces confident wrong answers
Process	Which workflows are documented, measured, and ready to redesign	Workflow redesign has the biggest effect on value ⁵

Dimension	What gets checked	Why it matters
People	Skills, training, readiness for change, and employee input into design	Adoption decides whether the gains appear
Security and privacy	Access control, data handling, and vendor and model risk	Protects customers and limits regulatory exposure
Governance	Decision rights, policies, risk review, and an inventory of AI in use	Decides whether AI can scale safely
Outcome measurement	A baseline and KPIs for each proposed use case	Without a baseline, value can't be proven
Readiness for AI callers	Whether your lines can serve customers' own AI agents that call, chat, or transact for them	Gartner expects a growing number of "machine customers" powered by AI agents ¹⁰

The four stages, from first steps to routine:



You come away with a scored report for each dimension, a list of use cases ranked by value and readiness, and the governance steps each use case needs before it goes live.

A Governance Model You Can Run

Governance answers four questions for every AI use case: who decides, what the rules are, how it's measured, and how often it's reviewed. The model below follows the four functions of the NIST AI Risk Management Framework (Govern, Map, Measure, and Manage).¹¹ It can be extended to meet ISO/IEC 42001, the first international standard for AI management systems.¹²

Decision rights

Executive sponsor. Sets AI strategy, budget, and how much risk the organization will accept, and answers for business outcomes and final escalations.

AI governance council. Operations, IT security, legal, HR, and a frontline representative. The council approves each use case and its guardrails, and owns policy, the AI inventory, and risk review.

Use-case owner. Runs one AI use case day to day and answers for its KPIs, incidents, and user feedback.

Data and knowledge owners. Decide what data and content AI may use, and keep it accurate, current, and limited to the people who need it.

Frontline champions. Help shape design, testing, and training, and flag problems early.

Before any use case goes live

1. It has a named owner, a written purpose, and a measured baseline.
2. Its access to data is limited to what it needs, and customer data is protected.
3. Customers are told they're dealing with AI wherever law or policy requires it.
4. There's a path to a person for exceptions, complaints, and high-stakes decisions.
5. It's tested for accuracy and unfair outcomes before launch, and on a schedule after.
6. Employees and customers have a way to report a problem, and every report gets a logged response.

Review schedule

Each use-case owner reviews KPIs against the baseline, incidents, and user feedback every month. The governance council meets quarterly to approve new use cases, update policy, review risk, and check the AI inventory. Once a year, or after a major platform change, the readiness assessment is repeated and priorities are reset.

From Assessment to Scale

The assessment is the first phase of the work. Each phase has to pass a gate, a measured result, before the next one starts and before more money goes in.

Five phases, each with a gate

The baselines set in phase 1 are what every later gate is measured against

- 1 Assess.** Inventory the data, score readiness, and rank use cases by value and readiness.

Gate: a baseline is measured for each use case

- 2 Govern.** The council approves decision rights, guardrails, and the AI inventory.

Gate: the use case is approved, with an owner and guardrails

- 3 Pilot.** Run one or two high-value use cases, designed with frontline employees.

Gate: the pilot beats its baseline on the agreed KPI

- 4 Scale.** Extend proven use cases across teams, with training before each rollout.

Gate: monthly KPIs stay on target

- 5 Sustain.** Monthly reviews, a quarterly council, and a yearly readiness re-score.

How Computer Instruments Approaches This

Computer Instruments has built self-service voice systems for public sector and enterprise organizations since 1979. Our team designs and integrates them, adding AI to the phone infrastructure an organization already has without disrupting it. For many organizations, the phone line is a sensible first use case. The routine questions are

easy to name, the call volumes are already counted, and the baseline is sitting in the call data.

The AI Concierge answers callers' routine questions in their own words, on the phone system you already run. It can tell callers they're speaking with AI when the greeting is set up to say so, and it can hand a caller to a person when they ask.¹³ That supports two of the guardrails in the checklist above.

If you'd like to start with the assessment, we'll give you a scored report for each dimension, a ranked list of use cases, and the governance steps each one needs before it goes live.

Sources

1. McKinsey & Company, "The state of AI in 2026: On the road to ROI," Dan Tinkoff, Lieven Van der Veken, and Michael Chui with Tara Balakrishnan, August 25, 2026 (online survey of 1,719 participants in 97 countries, May 4 to June 8, 2026).
2. Erik Brynjolfsson, Danielle Li, and Lindsey Raymond, "Generative AI at Work," *The Quarterly Journal of Economics* 140, no. 2 (2025): 889–942. Earlier version: NBER Working Paper 31161, April 2023.
3. Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," June 25, 2025.
4. Microsoft and LinkedIn, "Microsoft and LinkedIn release the 2024 Work Trend Index on the state of AI at work," May 8, 2024 (survey of 31,000 people in 31 countries).
5. McKinsey & Company, "The state of AI: How organizations are rewiring to capture value," March 12, 2025.
6. Perkins Coie (now Ashurst Perkins Coie), "Texas Enacts Responsible Artificial Intelligence Governance Act," July 3, 2025; Texas Legislature, "HB 149," 89th Legislature, signed June 22, 2025, effective January 1, 2026.
7. California Legislature, "AB 1609 (Zbur), Customer service chatbots," Chapter 733, Statutes of 2026, approved September 28, 2026.
8. Gartner, "Gartner Survey Finds 85% of Service and Support Leaders are Expanding Human Agent Responsibilities Despite Expectations of Mass AI Layoffs," April 28, 2026 (survey of 321 service leaders, September to October 2025).
9. Gartner, "Gartner Survey Finds 87% of Customers Say Companies Using GenAI for Customer Service Must Provide Access to a Human Agent," August 4, 2026 (survey of 3,566 customers, February to March 2026).
10. Gartner, "Gartner Predicts Agentic AI Will Autonomously Resolve 80% of Common Customer Service Issues Without Human Intervention by 2029," March 5, 2025.
11. National Institute of Standards and Technology, "AI Risk Management Framework" (AI RMF 1.0, NIST AI 100-1), released January 26, 2023.
12. International Organization for Standardization, "ISO/IEC 42001:2023, Information technology — Artificial intelligence — Management system," December 2023.
13. Computer Instruments product documentation for the AI Concierge: conversational answering on the phone system an organization already runs, an AI disclosure set in the greeting, and hand-off to a person on request.