humatica

November 2026

White paper

AI-Native by Design

The Operating Model Behind Lasting AI Advantage



In partnership with CoreWeave

Contents

Note from the authors 3

Summary 4

01 Leaders build organizations around AI, not AI around organizations 5

02 AI adoption is not AI maturity 10

03 Building got easier. Owning did not. 15

04 Enterprises are choosing the AI they can keep 19

05 The goal is to create tomorrow's business, not to automate yesterday's work 23

Five practices to turn AI ambition into lasting value 25

Methodology 27

About Humatica and CoreWeave 28

About this study

AI-Native by Design examines how enterprises organize, govern, and build with AI to create tomorrow's business. It draws on 30+ in-depth interviews with senior executives and private equity investors including leaders at Fortune 500 companies and Tier 1 private equity funds, and on a research panel of 252 enterprises across 14 industries from North America, Europe, and Asia, profiled between July and October 2026.

The study was conducted by Humatica in partnership with CoreWeave. The findings and views expressed are those of the authors.

© 2026 Humatica. All rights reserved.

Note from the authors

AI is changing how every business works, and each one now has to decide how to organize for it. Practice varies widely. Some enterprises have rebuilt roles, decision rights, and governance around AI, while others are still running pilots inside structures designed for a world without it. The underlying issue is that AI now does part of the work that people used to do: interpreting information, making decisions, and checking them. As a result, enterprises are reopening questions they settled years ago: how they are organized, how their operating model works, and the principles that govern who makes decisions, who checks them, and who is accountable for the outcome. These questions are the subject of this study.

AI enters the enterprise at three levels. It enhances individual work through assistants that help employees write, analyze, design, and code. It increasingly runs operations, making decisions inside processes that organizations depend on every day. And it is becoming part of the products and services companies sell, placing AI directly in front of customers. As AI moves from individual productivity to business-critical operations and customer-facing products, it becomes an organizational challenge as much as a technological one.

This shift creates new demands on governance. Access models designed for employees must now accommodate agents that can act faster, at greater scale, and across more information than any individual. Agents also need explicit rules on what they may decide alone and when they must escalate, judgments that managers used to make case by case. At the same time, AI is expanding who can build software. People outside traditional engineering functions can now create applications, workflows, and automations that others rely on. Governance must therefore extend beyond people to agents, and beyond IT to software built across the enterprise.

The challenge is not simply adopting AI but building the capability to run it. The same frontier models are available to almost every enterprise, so the models alone give no one an advantage. Advantage comes from systems built on the enterprise's own data and know-how, and from the teams and AI cloud platforms that turn models into production systems that improve through use. Because each system makes the next easier to build, early movers will be hard to catch. Deciding what to build, what to source, and where to run it is therefore a decision for the executive team.

To date, most of AI's value has come from efficiency. Those gains matter, but competitors using the same models will compete them away. The larger opportunity lies in work that was previously uneconomic, decisions that could not be made fast enough, and products and services that could not be built at all. AI lowers the cost of each. Because companies publicize their successes and rarely their failures, the gap between enterprises is hard to see from the outside, and we expect it to widen.

This study is written for executive teams whose businesses increasingly depend on AI and who are deciding how to organize for it. We owe our sincere thanks to the executives and investors who generously shared their time, experience, and candid perspectives.



Andy Knobloch

Partner, Humatica



Juan Zapardiel

Engagement Manager, Humatica

Acknowledgments

We thank our Humatica colleagues Charles Browning, Richard Breitschwerdt, Paul Simpson, Claudio Limacher, and Andros Payne, as well as our partners at CoreWeave for their valuable insights and contributions to this report.

© 2026 Humatica. All rights reserved.

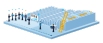
Summary

AI is rewriting what businesses can do to change the way they work and compete in the market. How much of that opportunity turns into value depends on the operating model. Companies that redesign ownership, decision rights, and structure for AI are building a lasting advantage.



01 Leaders build organizations around AI, not AI around organizations

An operating model for AI settles who decides what gets built, who owns each system once it runs, and who pays for it. Leaders make these decisions deliberately, move accountability for what gets built into the business, and redesign roles as AI takes on part of the work that existing organizations were built to manage.



02 AI adoption is not AI maturity

Counting users and use cases says little about whether a business can rely on the AI it has built. Maturity depends on what the business relies on AI to do and whether the organization can build, run, and improve those systems at scale. Closing the gap between adoption and maturity is now the main AI priority for executive teams.



03 Building got easier. Owning did not.

AI lets almost anyone in the business build a tool that others come to rely on, while the work of integrating, securing, monitoring, and paying for it has not shrunk. The risk sits with systems a team relies on and no one owns, and leaders reduce it by building these requirements into a shared platform.



04 Enterprises are choosing the AI they can keep

Most enterprises source AI capabilities externally for everyday work. The calculation changes for the workloads they compete on. There, leaders seek greater control over the systems, data, and knowledge that create advantage, so they can improve them over time and adapt them to their business.



05 The goal is to create tomorrow's business, not to automate yesterday's work

Most AI use cases make existing work cheaper or faster, and competitors using the same models can match those gains. Leaders also use AI to build capabilities they did not have before and products customers pay for, and we expect most of the value from AI to come from work of this kind.

Key insights

Enterprises are redesigning their operating models for AI and moving accountability for what gets built into the business.

One in five enterprises has named an AI leader, often with deep knowledge of the business.

Two-thirds of large enterprises now run a process, a product, or a business model that breaks if AI is switched off.

42% of AI use cases in production started in the business, and one in four was started by an individual or team on their own initiative.

73% of leaders name IP protection and differentiation as a reason for moving a critical workload onto infrastructure they control; 36% name cost and predictability.

Five practices to turn AI ambition into lasting value

- **A** Make AI an operating model decision
- **B** Empower decentralized building, and scale proven results
- **C** Unlock the enterprise's knowledge as a proprietary AI advantage
- **D** Keep control of the AI the business depends on
- **E** Keep every AI production system owned, measured, and improving

© 2026 Humatica. All rights reserved.

01

Leaders build organizations around AI, not AI around organizations



Most enterprises have not yet organized for AI. Technology now does part of the work that an organization's roles, handoffs, and approvals were designed around. Who decides what gets built, who owns each system once it runs, and who pays for it are often left open, and about a third of the enterprises in our panel have a formal operating model for AI. Enterprises typically organize AI in one of four ways, each trading control against speed of adoption. Accountability for what gets built is moving into the business. Leaders make these decisions deliberately and choose how to organize the people who build AI. Where enterprises have appointed an AI leader, we read the role as one of business transformation more than technology.

20%

of enterprises in our panel have appointed a formal AI leader

35%

of enterprises in our panel have a formal operating model for AI

© 2026 Humatica. All rights reserved.

Enterprises organize AI in four ways

An operating model sets how an organization groups its activities into roles, teams, and functions, who makes which decisions, and how the work is funded and governed.

Enterprises typically organize AI in one of four ways, each trading control against speed of adoption (Exhibit 1). Hub and spoke was the most common model among the enterprises we interviewed, often with a center of excellence retained for specific areas. Several had changed their model as their use of AI grew, and the leaders among them chose their model deliberately, knowing what it gave up.

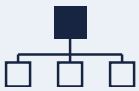
> *"Where we've landed is a CoE that works almost like a venture studio. Business units bring ideas into a funnel, we assess product-market fit, and a small cross-functional AI team spins them up fast. We expect most of them to fail, and that's fine. The point is to learn quickly and move on to the next idea."*
>
> Chief Data & AI Officer, Industrial Manufacturing

Exhibit 1

Enterprises organize AI in four ways, each trading control against speed of adoption

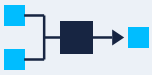
The four organizational models described by interviewees

01
Centralized center of excellence



A single central team owns all AI build and run. Best for high governance and tight risk control.

02
Centralized accelerator



Business units bring ideas; an accelerator builds with cross-functional teams. Best for new business models and building experience.

03
Hub and spoke



A central team plus function or business unit partners: central standards, federated execution. Best for balanced speed and governance.

04
Embedded in the business



AI developers in each team and no dedicated AI organization. Best for speed above standardization.



Control — **Speed of adoption**

> *"We never set up an AI steering committee. A lot of companies assembled a group to screen use cases, picked a handful, and then found each one was worth a few hundred thousand a year at most, so the POCs fail. Part of the reason we moved quickly is that we didn't have to consult others."*
>
> Chief Information & Operations Officer, Insurance

© 2026 Humatica. All rights reserved.

Accountability for AI is moving into the business

Enterprises are adding accountability for AI as well as capability (Exhibit 2). The roles that decide what gets built, and own the result, are moving into the business. Specialist teams keep responsibility for operating and securing what is built.

> *"Me driving adoption in a centralized way is actually now a bottleneck. Why do people need to go through me? So now we're changing our operating model."*
>
> Data & AI Engineering Executive, E-Commerce

Data scientists and machine learning engineers are no longer the only route to building with AI, and their capacity can be concentrated on the hardest problems, such as the AI that is itself the product. The CIO keeps the same mandate of security, identity, standards, and contracts, but now has to assure more systems, many of them built in the business.

> *"IT can no longer simply say no. It has to be useful, because otherwise people will go and do it anyway."*
>
> Chief Transformation Officer, Industrial Distribution

Exhibit 2

Accountability for deciding what AI to build is moving into the business, while specialist teams keep assurance

Roles involved in enterprise AI, grouped by what they are responsible for (data from panel and interviews)

Group	Role	Typical titles	Responsibility
Decide what to build	**Formal AI leader** New	Chief AI Officer, Head of AI	Sets priorities and allocates investment. The mandate varies across companies.
	Functional AI lead New	Owner inside functions like marketing, finance, or operations	Drives use cases for one function, usually alongside an existing role.
	Business owner	P&L holder, accountable for the outcome	Selects and funds use cases and owns the process changes and business transformation.
	Domain expert	Process owners, senior practitioners	Holds the process knowledge the system needs. Validates output and identifies edge cases.
Build it	**Citizen builder** New	Independent employee or team, enabled by the platform	Builds workflows and agents with low-code tools, without engineering support. Takes on the small, local automations and works inside guardrails set by IT.
	Data scientist or ML engineer	Applied ML inside product and engineering teams	Builds AI that is itself the product or the business model, such as physical AI, autonomy, and quantitative models. No longer the only route to building with AI.
Ensure it is built well	**AI platform owner** New	Head of AI or ML platform, VP of AI infrastructure	Owns the shared infrastructure: compute, model access, tooling, and monitoring.
	CIO or IT leader	Enterprise IT, security, identity, cloud contracts	The mandate has not changed significantly: security, identity, standards, and contracts. What has changed is the volume and origin of what has to be assured, since the business now builds for itself.

> *"The business decides what to build. IT ensures it is built well."*
>
> Commercial AI Lead, Industrial Manufacturing

© 2026 Humatica. All rights reserved.

The AI leader role is new, and its mandate is still forming

One in five enterprises in our panel (20%) has appointed a formal AI leader: a person who holds a mandate for AI across the enterprise, whatever their title (Exhibit 3). The role is most prevalent in sectors such as financial services. The title varies: 25% of these leaders carry a title specific to their company, 23% are Head of AI, and 15% are Chief AI Officer.

AI leaders promoted from within had spent considerably longer at the company before their appointment than CIOs had. Enterprises appear to give the role to people who know the business and how to change it, which we read as a role of business transformation more than technology.

> *"Companies are adopting AI leader roles now because running AI is not purely an IT job."*
>
> Chief AI Officer, Aerospace & Defense

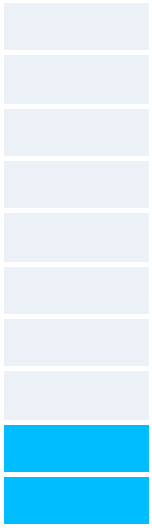
Exhibit 3

One in five enterprises has a formal AI leader, and those promoted from within had spent longer at the company than CIOs

Formal AI leaders in the enterprise panel, and their time at the company before appointment compared with CIOs

A Share of enterprises with a formal AI leader

Enterprise panel, n = 252

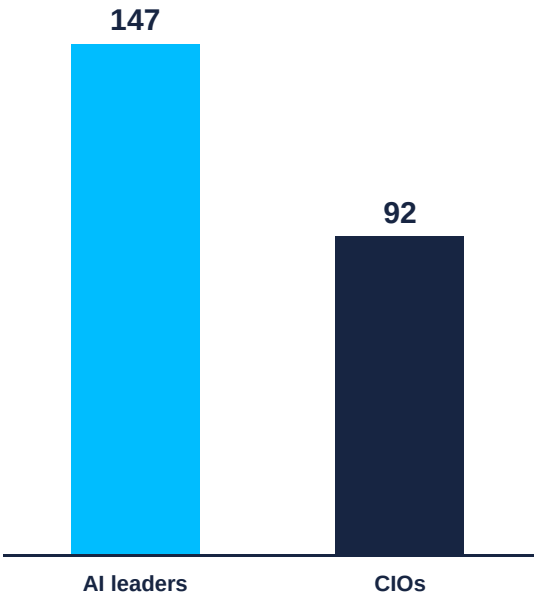


20%

of enterprises in our panel have appointed a formal AI leader

B Median months at the company before appointment

Internal appointments



Note: A formal AI leader holds an enterprise-wide mandate for AI, whatever the title. 44% of AI leaders and 42% of CIOs were internal appointments.

© 2026 Humatica. All rights reserved.

AI changes the costs an organization is built around

Much of an organization's structure exists to move information, route work to the people who can do it, and control error: status reports, handoffs, specialist teams, development queues, and approvals. Three costs that shaped this structure are falling.

Three costs that shaped organizational structure



01 The cost of

Coordination

Falls as AI tracks work and routes information directly to the people who need it, instead of filtering it upward through layers of management.



02 The cost of

Accessing expertise

Falls because the knowledge of experienced specialists, including some of their judgment, can be written down and made available to anyone in the business through AI.



03 The cost of

Execution

Falls as AI completes tasks that used to need a dedicated team: writing documentation, analyzing data, designing processes, and building software.

Each of these weakens the case for parts of the structure that exist today. Approval layers designed to catch human error are harder to justify where a model makes fewer errors and every step it takes is recorded. Roles that existed to pass information between teams lose part of their purpose, and work that waited in a specialist team's queue can be done by the people who need it. As a result, organizations will become flatter and work will be divided differently across roles. Senior employees take on broader responsibilities across domains, and some roles merge. For example, traditional product management and software development, which today separate deciding what to build from building it, can be combined in one person who takes a product from customer need to working release. Taken together, these changes alter how an organization is built.

> *"Organizations will get flatter, with senior people expanding horizontally across domains rather than managing upward hierarchies, so that information can flow directly to whoever needs it."*
>
> Chief Digital Officer, Private Equity Investor

© 2026 Humatica. All rights reserved.

02

AI adoption is not AI maturity



Adoption metrics such as number of active users, number of use cases, or AI spend do not show whether a business gets lasting value from the AI it uses, whether built in-house or sourced externally. We therefore assess maturity on two dimensions: what the business depends on AI to do (AI outcome), and whether the organization can build, run, and improve AI across the business (AI readiness). Enterprises need to build the organization at the same pace as their use of AI. Where readiness lags, the business relies on AI it is not yet able to improve or scale, and we see closing that gap as the most immediate opportunity.

95%

of enterprises in our panel have moved beyond individual experimentation

66%

run a process, product, or business model that would degrade or break if AI were switched off

12%

could not deliver their offer at its current price and cost base without AI

2.4x

as many interviewees named their operating model as a constraint on AI as named the technology

© 2026 Humatica. All rights reserved.

Maturity has two dimensions

Most measures of enterprise AI count adoption: how many employees use an assistant, how many use cases are live, how much is spent. Almost every enterprise in our panel uses AI in its daily work. However, these counts say little about whether the business has created more value, can rely on what it has built, or about what would happen if it stopped working.

We assess maturity on two dimensions (Exhibit 4). AI outcome describes the deepest way in which the business depends on AI, on a five-stage scale. Each stage is defined by what would happen if AI were switched off.

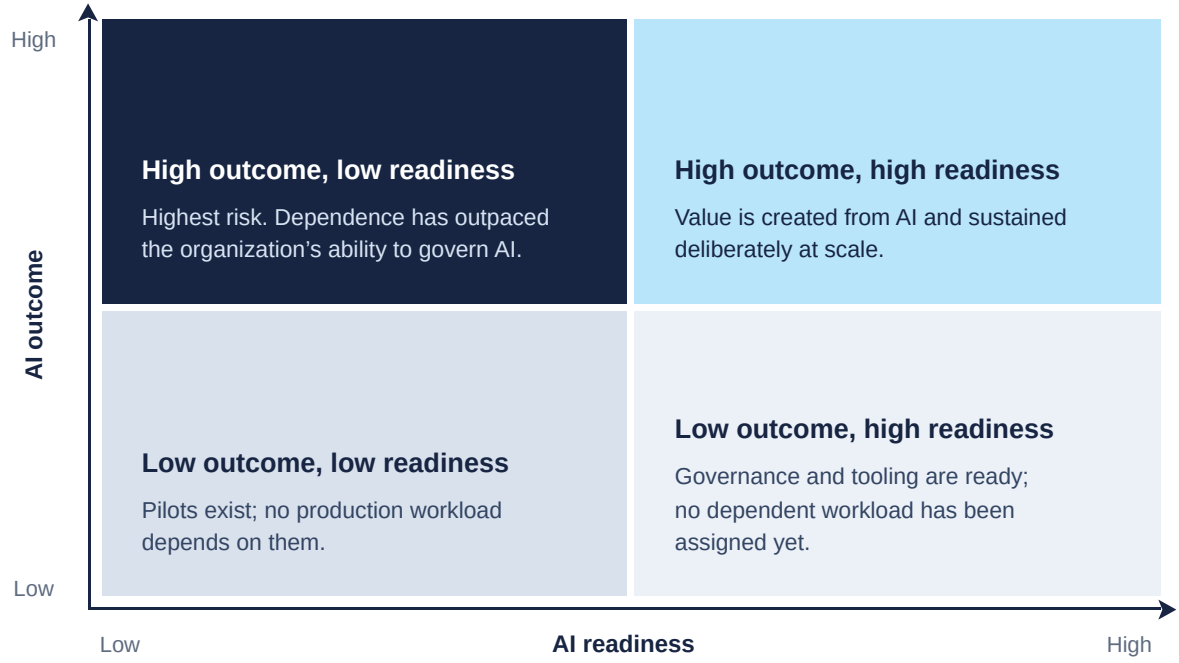
AI readiness describes whether the organization can own, govern, and extend that dependence. It covers five factors: organization and governance; guardrails, control, and assurance; roles and people management; data and proprietary knowledge; and technology and infrastructure. Four of the five are mainly organizational. In the panel, we use the maturity of each enterprise's AI operating model as the measure of readiness.

Exhibit 4

Illustrative

AI maturity combines what the business depends on AI to do with the organization's readiness to sustain it

The two dimensions of AI maturity, with the definitions used in this study



© 2026 Humatica. All rights reserved.

AI outcome stage	
1 Individual experimentation	Individuals use AI on their own initiative. No production process depends on it, and the value is limited to what one person can do.
2 Assisted work	AI is used widely in daily work with measurable gains, but processes do not depend on it. If the tool were removed tomorrow, nothing essential would break.
3 Process dependence	One or more core processes rely on AI to meet the required speed, cost, volume, or service level. Removing it degrades the business.
4 Product or service dependence	AI sits inside something the company sells, but pricing, revenue model, and unit economics are unchanged. Removing it breaks a product feature and some customer commitments.
5 Business model dependence	AI is in the economics of the business: it produces what is sold, prices it, or has replaced the labor that delivered it. Without it, the company cannot deliver its offer at its current price and cost base.

AI readiness factor	
☑ Organization and governance	The structure, decision rights, and operating model through which AI work is prioritized, funded, and run.
☑ Guardrails, control, and assurance	The testing, monitoring, escalation, and audit that keep an AI-dependent process within acceptable behavior, such as spend caps, approved model lists, evaluation, and the ability to switch a system off.
☑ Roles and people management	How roles, accountability, and performance management change once AI does part of the work, including review, approval, and escalation when output is contested.
☑ Data and proprietary knowledge	Whether the data a workload needs is accessible, documented, permitted, and current, and whether the enterprise's own knowledge is consistently curated.
☑ Technology and infrastructure	The platforms, tooling, and infrastructure available to build and run AI.

Note: Outcome stage is assigned to the deepest form of dependence an enterprise discloses. Readiness factors are those named by interviewees as constraints on progress (see Exhibit 6).

© 2026 Humatica. All rights reserved.

Two in three enterprises in our panel depend on AI

Of the enterprises in our panel, almost all have moved beyond individual experimentation, and two in three run a process, a product, or a business model that depends on AI (Exhibit 5). Switching AI off in these enterprises would degrade or break part of the business.

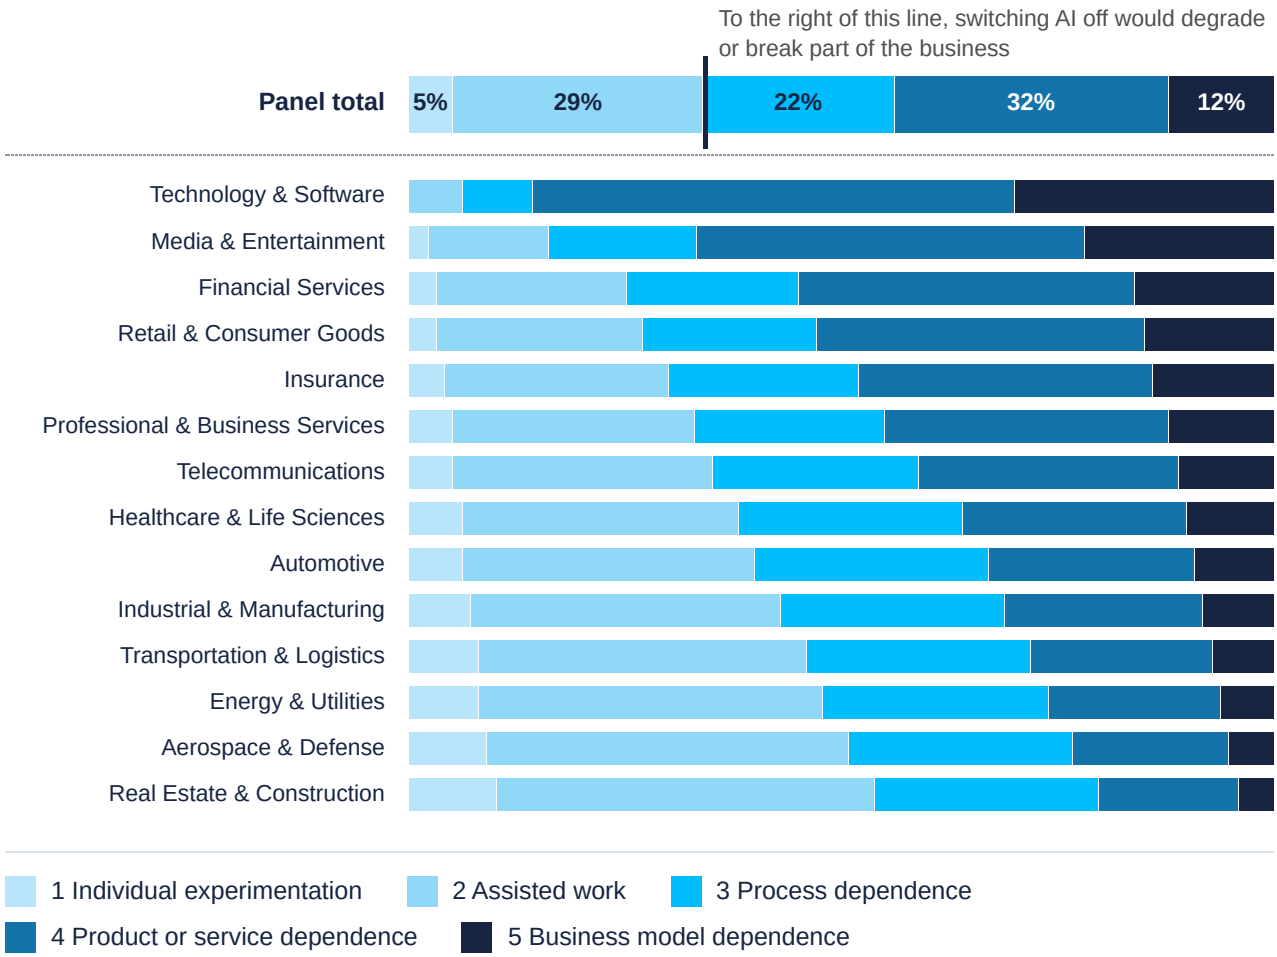
Adoption and dependence do not always go together. An enterprise can give every employee an assistant and still leave the work that runs the business untouched, or it can depend on AI in one critical process while most employees barely use it.

> *"We have 36 agents running the research funnel in R&D, finding unmet needs, doing market and IP landscape assessments. It is speeding up how we work, and it is producing actual good work, not generic slop."*
>
> President, Industrial Technologies

Exhibit 5

Two in three enterprises in our panel run a process, product, or business model that depends on AI

Share of enterprises at each AI outcome stage, panel total and by sector, %



© 2026 Humatica. All rights reserved.

Organization and governance hold back AI value creation

When we asked enterprise and investment leaders what limits their progress, they named their own organization as the constraint more than twice as often as the technology (Exhibit 6).

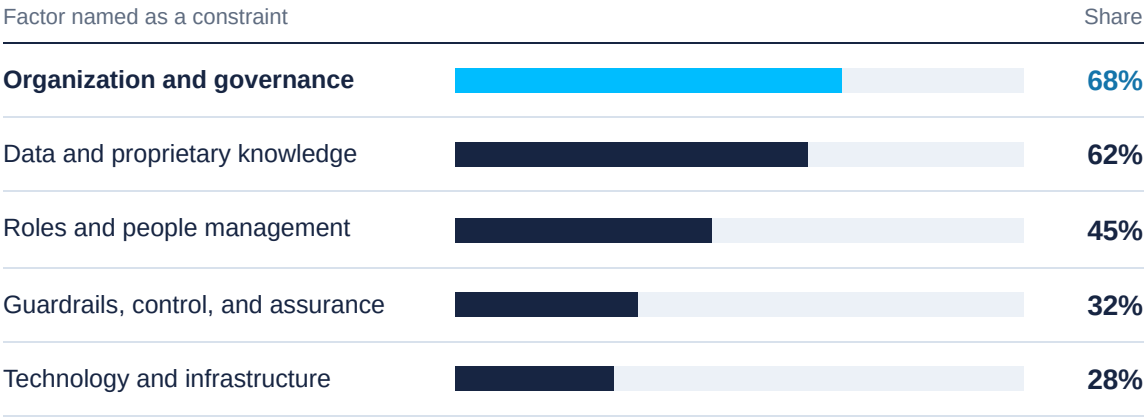
Several described the same sequence. AI spread quickly through the enterprise because the tools are easy to use, and early benefits were visible. The ownership, governance, and operating disciplines that an AI-dependent process requires came later, if at all. Interviewees found it harder to absorb AI into how the organization works than to acquire it: who approves what, who is accountable for the result, and how work moves between people and systems.

> *"This is primarily an organizational challenge, not a technology one. Existing approval processes and internal red tape are slowing progress and have become the main bottleneck."*
>
> AI Market Expert, China

Exhibit 6

Interviewees name the organization as the constraint more than twice as often as the technology

Share of enterprise and investment leaders who named each factor as a constraint on progress with AI, %



Note: Interviewees could name more than one factor. Responses were coded to the five readiness factors defined in Exhibit 4.

Readiness has not kept pace with dependence. Two in three enterprises in our panel depend on AI for a process, a product, or their business model, and about a third have a formal operating model for it. Many enterprises are therefore in the highest-risk position in Exhibit 4: they rely on systems that the organization is not yet set up to own. The same position holds the most immediate opportunity. These enterprises have already shown that AI works in their business, and building the organization around it is what lets them improve and scale it.

We think executive teams should make closing the gap between outcome and readiness their main AI priority. In practice this means knowing which processes and products now depend on AI, who owns each one, and whether the organization could detect and handle a failure in any of them.

© 2026 Humatica. All rights reserved.

03

Building got easier. Owning did not.



AI has made it possible for almost anyone in an enterprise to build systems that others come to rely on, and much of what now runs in production started in the business, often without a sponsor or a budget. The work of owning these systems has not shrunk. Someone still has to integrate, secure, monitor, pay for, and maintain each one, and AI adds tasks of its own, such as testing output that varies from one run to the next. Interviewees described this work as unowned in many of their enterprises. Leaders build the requirements of ownership into a shared platform, so that every new use case inherits them.

42%

of production AI use cases in our panel started in the business

24%

of production AI use cases were started by employees on their own initiative

© 2026 Humatica. All rights reserved.

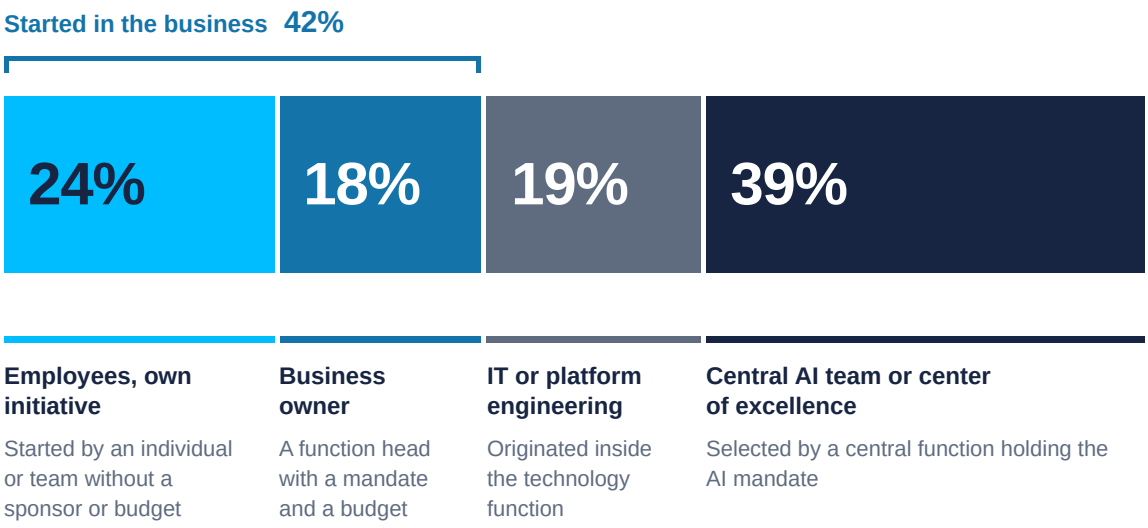
Anyone in the business can now build AI that others rely on

Until recently, building an application that a team could rely on took months of engineering work, and security, integration, monitoring, and support were settled before anything went live. Today a business user can build a working tool in days or hours. One engineering leader described setting up an AI assistant on a Saturday morning, specifying a service that had previously taken a team months to build, and having it running an hour later.

Exhibit 7

42% of production AI use cases started in the business, and a quarter were started by employees on their own initiative

Origin of production AI use cases, % of coded production use cases in the panel



Note: Production use cases are those an enterprise discloses as live, coded by where the use case originated.

Building is the smaller part of the work

> *AI can take a user from zero to capable very quickly. It also means a bad developer can now write bad code faster. Building is now only 30% of the work. Everything after that is data, integration, and architecture to operationalize it."*
>
> Chief Digital Officer, Executive Search

Running cost is now also harder to predict. Consumption pricing means cost grows with use, so a business case set during a trial can fail once the system is in production. The cost of AI can be significant and surprising: a product development unit saw its token bill skyrocketing without a corresponding step change in results. Businesses often discover that AI systems consume significantly more computing resources than expected once deployed at scale, eroding the economics of otherwise promising use cases. Without clear visibility into usage and cost, it becomes difficult to manage AI as a business asset.

© 2026 Humatica. All rights reserved.

Exhibit 8

Systems can now be built in days leveraging AI, but owning one takes more work than owning traditionally built software

What building and owning involve for systems built with AI, compared with traditionally built software

A Building: share of effort to build and launch an application

Schematic, illustrative shares



B Owning

Ten requirements must be considered alongside building

New
Did not apply to traditional software

Harder
Applied before, now more demanding

Unchanged
Same as for traditional software

Correctness	Output varies by run, so pass/fail tests cannot prove it
Model changes	Providers update and retire models on their own schedule
Prompt injection	Models cannot reliably tell instructions from data
Agent permissions	Agents act on systems and need permissions like a user's
Ownership	Systems built outside engineering often have no named owner
Visibility	Use cases go live without passing through IT
Skills	Few people can maintain an agent across its life
Data and audit	Boards and regulators want a record of what it did and why
Running cost	Cost grows with use and is hard to predict
Integration	Scaling up still means connecting to systems of record

Note: Requirements identified and classified by Humatica from interview evidence.

The risk sits with systems a team relies on and no one owns

The risk sits with systems that a user built and a team has come to rely on, since no one owns them (Exhibit 9). These systems appear without passing through IT, so the enterprise may not know how many it has.

> *Our voice-of-customer tool is driven by R&D, but who owns the business process it feeds is still unclear, and that needs work."*
>
> President, Industrial Technologies

> *We have 150 agents at the base and another 20 orchestrating them. My engineers can't oversee all of that on top of their day jobs. Someone has to own it, and jobs will have to change."*
>
> Chief Engineer, Automotive

© 2026 Humatica. All rights reserved.

By contrast, systems built by the platform team have a named owner. One industrial manufacturer has used agents to create sales orders in its ERP since early 2025; they now create more than 80% of orders, owned by the ICT function, with a human in the loop to check agent decisions.

> *This is a categorical problem at the moment. Teams build their own thing, and then they run their own thing, and then they leave, and the workflow falls over."*
>
> Data & AI Engineering Executive, E-Commerce

Exhibit 9

The risk sits with systems a team relies on that no one owns

AI systems by who built them and who relies on them (interviews)

Relied on by	Built by the platform team	Built by a user
A team or process	**Business-critical, owned** Example A customer-service automation, or an AI feature in the product Ownership A named owner, with escalation, change control, and monitoring	**Orphan service** Example RevOps builds a daily briefing sales now plans its day around Ownership **Nobody.** If it breaks, a team's process stops. No one to call
An individual	**Over-served** Does not occur in practice The platform team does not build for a single person	**Disposable** Example An analyst automates their own weekly report Ownership The user. If it breaks, they do the task by hand

Leaders build ownership into the platform

Without a shared platform, every team building AI must solve security, access, monitoring, and cost management on its own. One team told us that its engineers had been given tools to build agents, but not the tools to scale, monitor, or govern them. When multiple teams address the same operational challenges independently, the result is duplication, inconsistency, and unnecessary complexity.

Leading enterprises solve these problems once. They build security, access, monitoring, evaluation, and cost controls into a shared platform, so every use case built on it meets them from the start. The platform also shows the enterprise what is running and who built it, which is the first step to giving each system a team relies on a named owner.

> *Architecture has to be expressed as platform, so the standard way to build something is also the easiest. A new application can be running on our platform in 20 minutes with the security requirements already met."*
>
> Chief Information & Operations Officer, Insurance

© 2026 Humatica. All rights reserved.

04

Enterprises are choosing the AI they can keep



Most enterprises buy AI as a service for everyday productivity. The calculation changes for the workloads a business competes on. There, leaders want greater control over the systems, data, and intellectual property that create advantage and improve over time. They cite differentiation far more often than cost as the reason to bring these workloads under their control. Doing so also means taking responsibility for the data, evaluation, and operating capabilities needed to keep them improving.

73%

of the enterprise leaders we asked named differentiation and proprietary IP as a reason to move a critical workload onto infrastructure they control

36%

named cost as a reason to move a critical workload onto infrastructure they control

4–7 months

by which the best open-weight models released in mid-2026 trail the proprietary frontier

© 2026 Humatica. All rights reserved.

Control matters most where the business competes

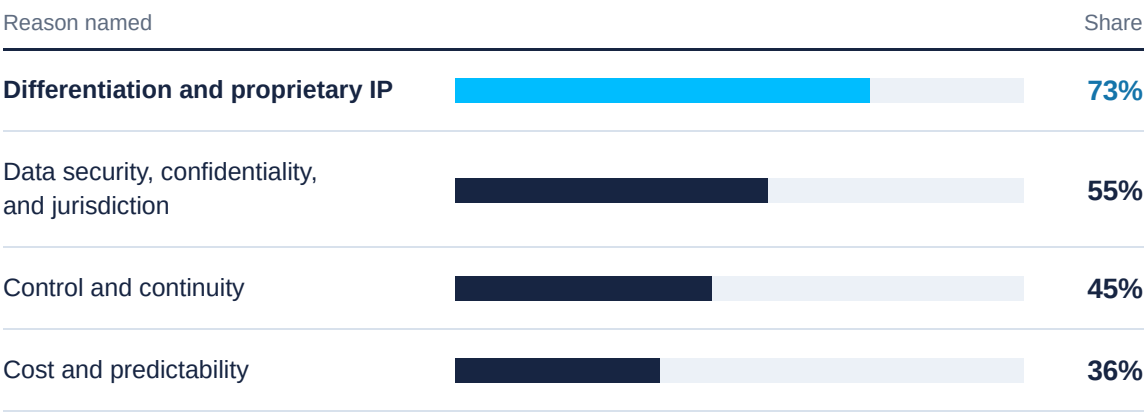
We asked enterprise and investment leaders what would lead them to move a critical workload onto infrastructure they control. They named differentiation and proprietary IP most often, and cost least often (Exhibit 10). Differentiation is a reason to own the model and the data.

> *Buy for commodity, build for differentiation. It is the same for AI."*
>
> Data & AI Engineering Executive, E-Commerce

Exhibit 10

Leaders most often cite differentiation and proprietary IP as a reason to move a critical workload onto infrastructure they control; cost comes last

Share of enterprise and investment leaders asked who named each reason for moving a critical workload onto infrastructure they control, %



Note: Interviewees could name more than one reason.

Two examples from our interviews show the kind of workload in question: a rental business matching customer requests against its stock, and an engineering team calibrating core systems with models built on its own test data. In both, the value comes from data that a competitor does not have.

> *One of our portfolio companies rents out more than 80,000 modular units, each configured differently. The tool finds the units closest to what a customer needs and easy to modify, so a salesperson can come back within the hour with delivery in two months rather than eight. The hard part was the data. It is still in test, with positive results."*
>
> Partner, Private Equity Investor

> *A hundred physical tests used to take us months. Now we run 15,000 simulations and only need a handful of physical tests to confirm the answer."*
>
> Chief Engineer, Automotive

© 2026 Humatica. All rights reserved.

No single model is best for every workload

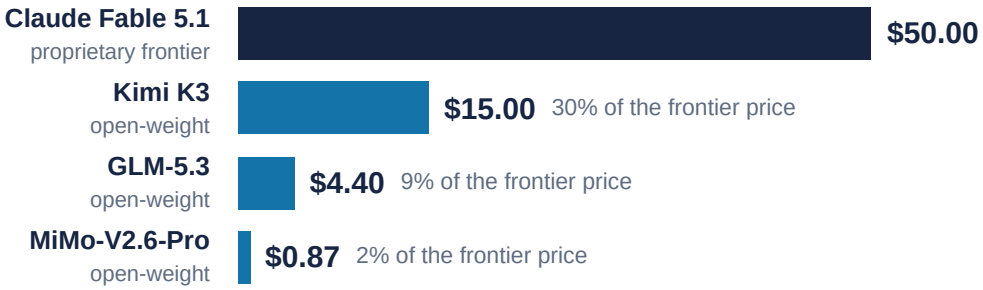
Several interviewees pointed to a recent cycle of open-weight releases as the moment their procurement conversations changed. Control has become a practical option as models evolve. On the Epoch Capabilities Index, the best open-weight models released in mid-2026 trail the best proprietary model by four to seven months, and leading open-weight models cost a third to a fiftieth as much as the frontier per output token (Exhibit 11).

For simpler tasks like retrieval, extraction, classification, and assistants inside products, the remaining differences in performance are no longer commercially significant. For the systems they build themselves, leading enterprises are increasingly combining open-weight and proprietary models, routing each task to a model according to its requirements, and reserving frontier models for complex tasks. The constraint some of them now meet is access to AI infrastructure.

Exhibit 11

Leading open-weight models cost a third to a fiftieth as much as the frontier per output token

List price per million output tokens, US$, October 2026



Source: Provider pricing pages

> *Building our own tools on top of commercial and open-weight models saves us at least $50 million a year on tokens. The bigger benefit is that we never had to ration access, because rationing brings fear back into an organization."*
>
> Chief Information & Operations Officer, Insurance

The advantage comes from improving AI with proprietary knowledge

As capable AI models become widely available, competitive advantage shifts from access to models to the ability to improve them. Leading enterprises combine frontier models with proprietary data, expertise, and operating know-how that competitors cannot buy.

They do so through a continuous improvement cycle. Every interaction, decision, and correction becomes a source of learning. Domain experts turn this experience into evaluations, training examples, and performance standards that reflect how the business operates. Models, prompts, and workflows are then refined against those standards, with each new version tested before deployment. Over time, the system becomes increasingly tailored to the enterprise and harder for competitors to replicate.

> *The most advanced portfolio companies are moving to open-source models and training their own, with routing across models on top."*
>
> Chief Digital Officer, Private Equity Investor

The question for executives is how to acquire the capability required to build and sustain these systems.

© 2026 Humatica. All rights reserved.

Five ways to close the capability gap

Most enterprises already have the ingredients that create advantage: the understanding of the problem they want to solve, the domain expertise, and the data. What they often lack is what execution requires: machine learning and platform engineering, evaluation, and the ability to run AI at scale. Leaders have multiple options to tackle this challenge, summarized as five options below. The right choice depends less on technology than on how much control the enterprise wants to retain over the AI capabilities that matter most (Exhibit 12).

Leaders decide use case by use case, weighing the questions in Exhibit 12. Three conditions put a workload on capacity the enterprise controls, regardless of the other answers. The first is that the data cannot leave the enterprise's perimeter because of regulation, contracts, or sovereignty rules. The second is that the workload must respond within a fixed time or run physically close to the site it serves. The third is that the serving stack itself has to be changed in ways a managed endpoint does not allow.

> *"We have not found the right people. The AI experts lack the domain expertise, and the people who know the domain lack the AI skills."*
>
> Operating Partner, Private Equity Investor

Exhibit 12

Enterprises have five ways to close the gap between AI ideas and execution, and the option to defer; each keeps a different degree of control

Options for acquiring the capability to build and run AI, from most to least control, and the questions to weigh for each use case

A Options

Control		Option	What it keeps	Typically suits
■■■■■	1	**Hire and develop an in-house team**	The capability, the model, and the IP	Use cases the business competes on, at sustained scale
■■■■□	2	**Forward-deployed engineers from technology providers**	The model, data and IP, and skills if transfer is planned	Differentiating work that must move faster than hiring allows
■■■□□	3	**Engage a systems integrator**	The solution; much of the know-how leaves with the integrator	Large programs that must integrate with existing systems
■■□□□	4	**Partner with a vertical specialist**	Use of the specialist's model, with data rights set by contract	Domains where the specialist's model and data are ahead of the enterprise's own
■□□□□	5	**Outsource the process**	The outcome; the capability stays with the provider	Non-core processes where buying a result is enough
Or		**Defer, and revisit when the technology has advanced**	Cash and options, at the cost of time while others build	Non-critical work while the technology or economics are still moving

B Dimensions — Questions to weigh for each use case

Strategic control
Is this where the business competes?

Domain specificity
How much does it rely on in-house data and know-how?

Criticality
What would a failure cost the business?

Scale
How much volume, and for how long?

Business case
Does the value justify the build?

Timing
How urgently is it needed, and how fast must it be implemented?

© 2026 Humatica. All rights reserved.

05

The goal is to create tomorrow's business, not to automate yesterday's work



Most of the AI use cases we see make today's work cheaper or faster. Those gains are becoming real, but competitors using the same models can match them. Most of the value from AI will come from what it makes possible, in two places: inside the business, where processes can do what they could not do before, and in what the company sells, where AI creates new products and services that customers pay for.

> *We have all the data you could imagine. The only limit now is what I can think to do with it, and whether we can execute."*
>
> Chief Transformation Officer, Industrial Distribution

Efficiency gains are open to every competitor using the same models, and we expect them to be competed away. The examples of new capability we found were specific to each business: an engineering team that now runs 15,000 simulations in place of a hundred physical tests, and a diversified industrial company running 36 agents across its research and development funnel.

> *I don't want to reduce my team just to hand the savings over to an AI-model company. The real opportunity is doing things we simply couldn't do before."*
>
> Chief Marketing Officer, E-Commerce

© 2026 Humatica. All rights reserved.

Exhibit 13 illustrates how leading enterprises are applying AI across business functions at different levels: individual productivity, processes, and products and business models. The examples are based on our enterprise panel, executive interviews, and Humatica project cases. They are illustrative rather than exhaustive, and every function shows opportunities across all three levels.

Exhibit 13

Leaders use AI in every function, from individual productivity to the products they sell

Examples of how leaders use AI, by function and by where the AI sits in the business (non-exhaustive)

General	Sales and marketing	Customer service and aftersales	Operations and supply chain	Product and R&D	Corporate functions
Products and business models AI in what the company sells					
Physical AI: autonomous field equipment	**AI search APIs**	**Premium AI service companion**	**Supply network visibility products**	**Code builder platforms**	**AI buying agents**
Clinical decision support systems	**AI sales configurator offered to customers**	**Enterprise customer service agent platforms**	**Supply chain intelligence as a service**	**Autonomous design tools**	**AI legal platforms**
Drug discovery platforms	**AI pricing advisor**	**Predictive maintenance offerings**	**Inventory optimization subscriptions**	**Material science platforms**	**AI financial analysis and diligence platforms**
Process: new capabilities AI lets a process do what it could not do before					
Enterprise knowledge search	**Cross-sell and next-best-offer engine**	**Predictive customer intervention**	**Computer-vision quality inspection**	**Virtual prototyping**	**Continuous forecasting**
Real-time multilingual collaboration	**Dynamic pricing intelligence**	**Predictive service renewal**	**Supplier risk and lead-time prediction**	**Generative product design**	**Continuous control monitoring**
Self-service data analysis for non-analysts	**Predictive opportunity identification**	**24/7 multilingual service support**	**Real-time demand sensing**	**Large-scale engineering simulation**	**AI-assisted regulatory change monitoring**
Process: tasks automated AI takes over tasks in an existing process					
Note taking	CRM updates	Customer support / tier-1 agent	Manual visual inspection	Documentation maintenance	Order-to-cash: credit risk and collections
Transcription	Lead enrichment	IVR replaced by voice agents	Purchase order matching	Manual test scripting	Zero-touch invoice processing
Document formatting	Contact data maintenance	Service scheduling, dispatch, and triage	Routine operational reporting	Engineering file organization	User provisioning
Individual productivity AI helps people with their own work					
Research assistant	Proposal / RFP companion	Service technician companion	Production planning support	Coding copilot	Business performance analytics
Content drafting (emails, docs, slides)	Campaign creation	Automated service reporting	Spare parts demand forecasting	Requirements generation	Recruiting screening
Meeting preparation	Sales call preparation copilot	Case summarization and recommendations	Technician copilot for manuals and troubleshooting	Test creation and automation	Coding assistants and agents

Note: Illustrative and non-exhaustive. Examples are drawn from enterprises in our panel, our interviews, and Humatica's client experience.

© 2026 Humatica. All rights reserved.

Five practices to turn AI ambition into lasting value

Humatica has identified five practices that will help executive teams turn AI ambition into lasting value. They draw on what leaders told us about their own businesses, the data in this report, and Humatica's work with clients. Most begin with a decision that only the executive team can make.

A Make AI an operating model decision



AI now completes some work faster than organizations can govern it. Approval chains, handoffs, and roles designed around human work have largely remained unchanged, even as AI takes on more of that work. The challenge is no longer simply automation, but redesigning how decisions are made and accountability is assigned.

Leading enterprises treat these as decisions for the executive team. The executive team sets the operating model for AI, names an owner for each process and product that depends on it, and ties AI investment to the budget holders who answer for the result. It also decides what the enterprise will build in-house, what it will buy, and what it will source from partners. Most enterprises hold the domain expertise that AI needs, and few will build large machine learning engineering or research teams of their own, so this decision determines how much of the capability the enterprise keeps.

> *"A poorly designed process that runs faster on AI can now create more issues faster than before, instead of creating better impact."*
>
> Chief Data & AI Officer, Industrial Manufacturing

B Empower decentralized building, and scale proven results



Building is no longer confined to engineering: 42% of the production use cases in our panel started in the business. Leaders who tried to route all development through a central team found that the team became the bottleneck, and the business built around it anyway.

Leading enterprises give the business a shared way to build. Shared tooling, common data connections, reusable components, and standard security controls allow many teams to build without each having to solve the same problems independently. Anything that goes into production passes a stricter gate: it is registered, it has a named owner, and it meets the security and evaluation standards set for systems that others depend on. Most enterprises generate AI value through a portfolio of use cases rather than a single breakthrough application. Few of those use cases justify bespoke development. When use cases are built on shared components, such as data connections, evaluation sets, and reusable agents, each one makes the next cheaper and faster to deliver. The enterprise can then let many ideas start and put its investment behind those that prove their value in use.

> *"If you want to stifle innovation, mandate exactly how people should do something. You still need a central place to share what works. The real test comes after you've proven an idea: can you scale it?"*
>
> Head of AI Strategy, IT Services

© 2026 Humatica. All rights reserved.

C Unlock the enterprise's knowledge as a proprietary AI advantage

General-purpose models give every competitor the same capability at the same time, and the efficiency it delivers can be matched by rivals or priced away.

The advantage lies in data, expertise, and context that competitors cannot replicate. Judgment that was never written down can now be codified, and together with the enterprise's own data they form sustainable IP.

Knowledge creates an advantage only when it is curated. Leaders give it an owner and a quality standard and keep experts in the loop to maintain it and to check what models do with it. Where a workload carries competitive consequence, they build or tune models on this knowledge to create differentiation that competitors cannot copy.

Leaders are also selective. Making everything machine-readable carries cost and risk, and not all knowledge is worth codifying. They start with the knowledge that matters most to how the business competes.

> *"You need to own the data, to own the outputs, to own the models that run on top of that, or to control them. Everybody can buy the pieces if you have the money. The question is how you integrate and stay one step ahead."*
>
> Chief Marketing Officer, E-Commerce

D Keep control of the AI the business depends on

AI dependence is becoming an operational risk. As AI becomes embedded in critical processes, products, and business models, leaders are preparing for scenarios in which providers change pricing, alter terms, restrict access, run short of capacity, or withdraw services altogether.

Leaders design flexibility into the AI architecture. Work is routed across multiple models based on task requirements, cost, performance, and sensitivity. Frontier models are reserved for high-impact use cases, while sensitive workloads remain on infrastructure the enterprise controls.

Leaders treat compute as a dependency in the same way. They secure capacity for the systems the business relies on, and run their most demanding AI workloads on infrastructure built for AI.

Strategic advantage comes from combining different models with proprietary knowledge. Together they enable both breakthrough innovations and unique products, services, and processes.

> *"We are not betting on one player, and we need to avoid vendor lock-in in all parts."*
>
> Chief Technology and Operations Officer, Traffic Technology

E Keep every AI production system owned, measured, and improving

Cheap building produces more agents than anyone maintains. One in four production use cases was started by employees on their own initiative, and a workflow or agent often stops working once the person who built it leaves.

The value comes from what happens after the launch. Each system has an owner who tracks its results and people in the business who correct its output, and on a regular cycle those corrections are used to improve it, or the system is retired.

Systems improve through use and through better models. Leaders keep a record of what each system does and what it costs, and set evaluations that any new model must pass before release. They design systems to take a better model without being rebuilt, since models improve every few months, and make sure the engineering and evaluation skills to do this are available, built in-house or brought in from AI specialists.

> *"If you don't keep improving your agents once they're in production, and you don't put proper governance around them, you end up with an agent wasteland."*
>
> Chief Engineer, Automotive

© 2026 Humatica. All rights reserved.

Methodology

The study combines more than 30 in-depth interviews with senior executives and investors, including leaders at Fortune 500 companies and leading private equity funds, with a panel of 252 enterprises. Both were carried out between July and October 2026. This study was conducted by Humatica in partnership with CoreWeave.

30+
in-depth interviews

252
enterprises

14
industries

Interviews

Interviewees include chief executives and business unit leaders, chief AI, data, and digital officers, CIOs, partners at private equity and infrastructure investors, and academic experts. Each interview followed a common core of questions on AI budget ownership, authority over deployment, changes to organization and roles, and sustaining AI in production, with further questions specific to the interviewee's industry and role. Responses were coded against the common core and compared across industries and ownership types. Interviews were conducted on a non-attributable basis. Quotations are attributed by role and industry and have only been edited for length and clarity.

Research panel

The panel covers 252 listed, private equity-backed, and family-owned enterprises with active AI programs, across 14 industries. AI labs, AI and cloud vendors, startups, and companies with revenue below $50 million were excluded. Each enterprise was profiled from primary disclosure, annual reports, regulatory filings, and investor materials, drawing on more than 1,900 documents and 6,900 sourced citations. Enterprises were placed on a five-stage scale of AI dependence, from individual experimentation to business model dependence, and assessed on the maturity of their AI operating model. Production use cases were coded by origin and by the outcome they deliver, and AI leadership roles by title and by time at the company before appointment. All coding was independently reviewed.

Authors

Andy Knobloch
Partner, Humatica · andreas.knobloch@humatica.com

Juan Zapardiel
Engagement Manager, Humatica · juan.zapardiel@humatica.com

© 2026 Humatica. All rights reserved.

About Humatica

Humatica is a specialist advisor on operating model transformation to private equity funds and enterprises. It combines proprietary organizational benchmarks and analytics with hands-on support, helping leaders build organizations that deliver value and growth, including designing AI operating models that turn ambition into lasting competitive advantage. Humatica has offices in London, New York, Munich, Paris, Zurich, and Singapore. Learn more at www.humatica.com.

About CoreWeave

CoreWeave is The Essential Cloud for AI™. Built for pioneers by pioneers, CoreWeave delivers a platform of technology, tools, and teams that enables innovators to move at the pace of innovation, building and scaling AI with confidence. Trusted by leading AI labs, startups, and global enterprises, CoreWeave serves as a force multiplier by combining superior infrastructure performance with deep technical expertise to accelerate breakthroughs. Established in 2017, CoreWeave completed its public listing on Nasdaq (CRWV) in March 2025. Learn more at www.coreweave.com.

This report is based on interviews and on information that enterprises have disclosed publicly. It is provided for general information and does not constitute investment, legal, or other professional advice.

© 2026 Humatica. All rights reserved.



AI-Native by Design

The Operating Model Behind Lasting AI Advantage

In partnership with

humatica CoreWeave

humatica.com

© 2026 Humatica. All rights reserved.