N-iX

Executive brochure

AI ambition is running ahead of enterprise data

The N-iX Data Index 2027: Closing the AI-readiness gap

Most enterprises now have an AI mandate, an AI budget and a modern data platform. **What most lack is data that AI can use.** The constraint has moved from the model and the budget to the state of the data, and to the lack of anyone accountable for it.



80%

of engagements began with the client wanting more from AI than the data could support



75%

of clients had documented a third or less of the critical data problems the engagement found



1 in 185+

enterprise data requests asked for AI-ready data by name

WHAT THE INDEX IS BUILT ON

185+

enterprise data opportunities from the N-iX pipeline

258

data and AI practitioners surveyed on their latest engagement

14

audience polls

Leader interviews

Data and AI leaders at enterprises including American Express, PayPal, Meta and Vertex Pharmaceuticals

Why the bill has come due now

A BI strategy was written for an analyst. The analyst absorbs a late pipeline, a manual correction or an undocumented definition before anyone else sees it. Models and agents consume data directly and continuously, so every deferred decision about quality, ownership and documentation is now exposed in production.

89% of data leaders whose companies run AI in production have seen inaccurate or misleading outputs (Salesforce)

What we mean by AI-ready data

Data that is:

1. Aligned to a named use case
2. Governed at asset level
3. Quality-checked in the pipeline
4. Described by live metadata and lineage
5. In scope across structured and unstructured formats

What engineers found when they opened the data

Source for all charts: N-iX Data and AI Engineering Survey, September 2026.

65%

A modern platform does not close the gap

Engineers rated 40% of the platforms they walked into as modern and AI-ready. About two-thirds of those still could not support the AI the client wanted. On modern platforms built only for BI, all but one engagement showed the gap.



82%

The failure is semantic

Undocumented business logic (54%) and conflicting sources of truth (28%) make up 82% of the quality failures engineers meet most often. Schema drift, the failure most quality tooling targets, makes up 5%.

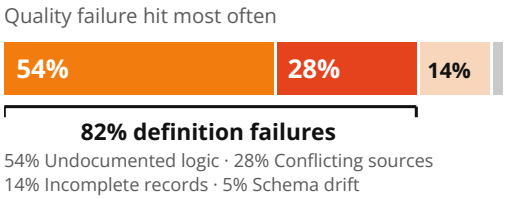

63%

Nobody owns it

In 63% of engagements no one structurally owned making data AI-ready. The work fell to the AI project team or to no one. 84% of engineers hit a significant blocker in the first 90 days, and most of those blockers were organizational.

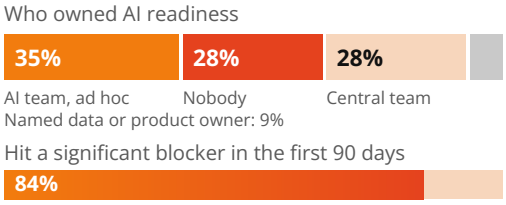

61%

Evaluation is the unsolved part of GenAI

Among engineers who have prepared unstructured data for GenAI, 61% name evaluation as the hardest part: knowing whether retrieval is good enough to act on.

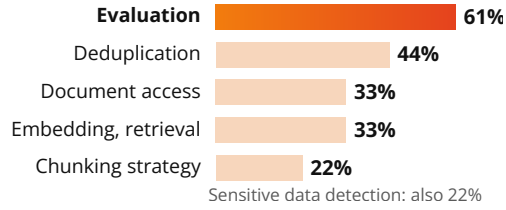

1 in 7

The ready minority planned for foundation work

14% of engagements began with ambition matched by data. These clients still did foundation work. They had named data owners, written definitions, a narrow first use case and a timeline that included preparation. None of it required a new platform.



What it means for leaders

The AI investment with the best return sits upstream of the model: naming owners, agreeing definitions, putting quality gates in pipelines, and bringing the unstructured majority of the estate into scope, all sequenced against one narrow use case. Organizations that did this spent their AI budget on AI. The rest paid for the same foundation work inside every pilot and recorded it as a delay.

Five dimensions, five owners

Quality and metadata respond to focused engineering within a quarter. Architecture and operating model take several. A readiness profile is more useful than a single score.

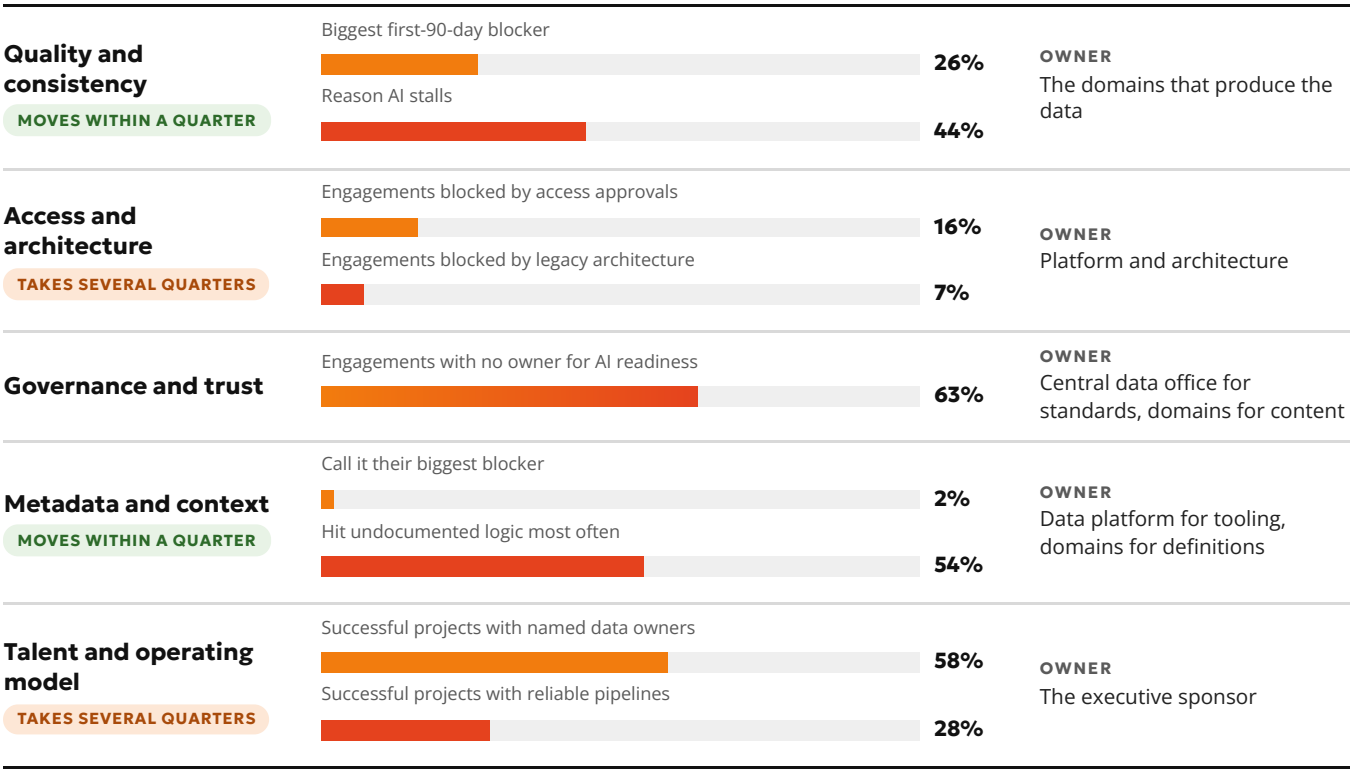

Source: N-iX Data and AI Engineering Survey, September 2026. Bars run from 0 to 100%.

Four archetypes, four first moves

The Sprinter

An AI mandate from the top, pilots in several business units, no owner for the data under any of them.

FIRST MOVE
Name an owner, pick one use case, pause parallel pilots.

The BI-Locked

A modern platform built for dashboards, with a roadmap that treats the model as a smarter dashboard.

FIRST MOVE
Audit the estate against the AI-ready definition, not the BI one.

The Gated

A sound platform and governance built for quarterly reporting, with every AI use case waiting in a review queue.

FIRST MOVE
Scope governance per use case and convert gates to guardrails.

The Ready Minority

Ambition matched by data, with foundation work planned, owned and funded.

FIRST MOVE
Institutionalize, then move from one use case to a portfolio.

> The archetypes describe where an organization is today, not stages it moves through in order. The distance to readiness is organizational before it is technical.

One objective: a single AI use case in production

The use case runs on data that meets its stated readiness conditions, and the preparation is visible and reusable. Six moves, in order:

1. **Name the owner**
2. **Pick one narrow use case**
3. **Write the definitions**
4. **Put quality gates on the pipeline**
5. **Hold a one-day governance session**
6. **Ship and measure**



A gate

Stops every use case until a committee reviews it.



A guardrail

Defines in advance what is permitted for a class of data and a class of use, and lets anything inside the boundary proceed.

1 day

With legal, security, data and business owners in one room, replacing months of sequential approvals. Approvals were the biggest early blocker in one engagement in six.

From AI-ready to agent-ready

Agents retrieve data across systems, chain decisions and act, with no human in the loop to catch exceptions. That raises the bar on semantic context, on currency, and on governance enforced at machine speed.

The decisions that make an estate AI-ready in 2027 are the same ones that give the agents of 2028 somewhere to stand.

The readiness gap for agents

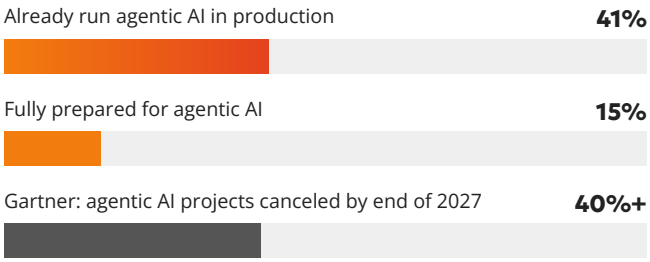

Sources: Fivetran (41%, 15%); Gartner (over 40%).

Get the full report →