

Beyond the AI Pilot

A Framework for Building Systems That Actually Deliver

Kanishk Mittal

AI Council • May 26, 2026

Show of Hands

“How many of you have had an AI pilot that never made it to production?”

“How many have had AI give confidently wrong answers because it didn’t have the right context?”

“How many have had AI recommend something - but a human still had to manually execute every step?”

95%

of enterprise generative AI pilots fail to reach production.

Let's talk about why this keeps happening.

It's Not the Models

What people blame

- Model accuracy & benchmarks
- Training data quality & coverage
- Vendor & platform selection
- Team skill gaps & hiring

What actually kills pilots

- No live business context
- Can't act across enterprise systems
- No path from vision to execution

Two Gaps, One Problem

The Infrastructure Gap

X Can't Read

AI has no governed access to live business data

X Can't Write

AI has no safe pathway to take action in systems

The Execution Gap

X No Execution Discipline

- Vision never becomes a project
- Projects never become tasks
- Nobody owns "what happens this week"

Three tools to close both gaps: a Read Contract, a Write Contract, and an Execution Ladder.

The Read Contract

Giving AI governed, auditable access to live business context.

What Is a Read Contract?

A formal agreement defining how AI accesses your business data: what sources it can read, how fresh that data must be, what governance applies, and what happens when data is unavailable.

Sources

Which systems does AI have access to?

Freshness

Real-time, daily, or stale?

Governance

Who authorized this?
Is it auditable?

Fallback

What does AI do when data is unavailable?

Name the Failure

A company deployed a customer-facing chatbot. A customer asked about a bereavement fare discount. The chatbot confidently explained a refund policy that did not exist.

The customer booked based on this advice. The company was taken to court - and lost.

What went wrong?

Air Canada: Flying Blind

- Chatbot fabricated a bereavement fare refund policy
- Customer relied on the fabricated advice and booked
- BC Civil Resolution Tribunal ruled against Air Canada (Feb 2024)
- Company held legally liable for its chatbot's statements

Read Contract Failure: No access to authoritative policy source

Name the Failure

A company used AI algorithms to buy and flip homes at scale. The AI priced offers based on historical data and automated valuations.

The algorithm consistently overpaid. The company wrote down \$500M+ and shut the division entirely. 2,000 employees lost their jobs.

What went wrong?

Zillow Offers: \$500M Read Contract Failure

- AI-powered home purchasing using the Zestimate algorithm
- Read historical sales - but not real-time local signals
- Phoenix: 90% of purchased homes listed at a loss
- \$500M+ write-down, 25% workforce layoff, program shutdown

Read Contract Failure: Historical data without real-time context

What a Good Read Contract Looks Like

Without a Read Contract

- Siloed, stale data
- No audit trail
- AI hallucinates when data is missing

With a Read Contract

- Unified real-time view
- Governed, auditable access
- Graceful degradation

JPMorgan COiN: Direct access to 12,000 commercial loan documents. 360,000 lawyer-hours replaced. 80% error reduction.

Works because the AI reads a specific, structured, authoritative source.

The Write Contract

Giving AI safe, auditable pathways to take action.

What Is a Write Contract?

A formal agreement defining how AI executes actions in your systems: what actions it can take, under what constraints, with what approval workflows, and with what audit trail.

Scope

What actions can AI perform? What is off-limits?

Constraints

Dollar limits, escalation thresholds, time windows

Approval

Autonomous, human-in-the-loop, or human-on-the-loop?

Audit

Every action logged, traceable, reversible

Name the Failure

A company deployed an automated trading system that could execute stock market transactions without human intervention. In 45 minutes, it executed \$7 billion in errant trades.

There was no dollar cap, no kill switch, and no escalation threshold. The company lost \$440 million in a single morning and nearly went bankrupt overnight.

What went wrong?

Knight Capital: \$7B in 45 Minutes

- Automated trading system deployed with untested code change
- No dollar cap, no position limit, no kill switch
- \$7B in errant trades executed in 45 minutes
- Company lost \$440M, rescued via emergency equity sale

Write Contract Failure: No scope limits, no constraints, no escalation threshold, no reversibility

What a Good Write Contract Looks Like

Without a Write Contract

- AI recommends but cannot execute
- No defined scope or constraints
- No audit trail when things go wrong

With a Write Contract

- AI executes within defined scope
- Tiered autonomy with guardrails
- Every action logged and reversible

Klarna AI Assistant: Handles 66% of customer service chats. Defined scope (refunds, order changes), hard limits (no account deletion), full audit log.

Works because the AI acts within explicit boundaries - and every action is traceable.

The Execution Ladder

Why architecture alone is not enough.

The Execution Gap

42% of firms abandoned AI initiatives in 2025, up from 17% in 2024. These are not technical failures. They are organizations that could not turn "do AI" into something a team can execute.

S&P Global Market Intelligence, 2025

Board says

"We need an AI strategy"

Team hears

"Build something with AI"

Six months later

Disconnected pilots, no measurable outcomes, no path to production

Name the Failure

A company invested billions in AI to "transform healthcare." They acquired multiple AI startups, hired thousands of engineers, and announced partnerships with major hospitals.

Five years later, they sold the division at a massive loss. The AI initiative was never decomposed into specific, measurable clinical workflows - it remained a vision without sub-projects or next actions.

What went wrong?

IBM Watson Health: Vision Without Execution

- \$4B+ invested across acquisitions (Truven, Phytel, Merge Healthcare)
- Vision: "Transform healthcare with AI" - no measurable outcomes defined
- No decomposition into scoped projects or weekly deliverables
- Sold to Francisco Partners in 2022 at a fraction of investment

Execution Ladder Failure: Vision without decomposition - empty middle of the ladder

The Execution Ladder

Vision

In 3 years, 60% of tier-1 support resolved by AI.

Outcomes

Reduce resolution time by 40% by Q4 2027.

Domains

Data infrastructure, API readiness, governance.

Projects

Build Read Contract for CRM + ticketing system.

Next Actions

Map all data sources. Identify API gaps. This week.

Diagnostic: Score Yourself (1-5)

Read Contract

Real-time data access?

Governed and auditable?

Graceful degradation?

Write Contract

API-driven actions (not UI)?

Defined autonomy tiers?

Every action logged?

Execution Ladder

Tied to measurable outcome?

Decomposed into projects?

Someone owns this week?

90 seconds. Use your worksheet. Where are you strong? Where are you stuck?

Take Out Your Phone

1. Audit your Read Contract

Map every data source your AI touches. For each: real-time or stale? Governed? What happens when it is down?

2. Draft your Write Contract

List every action your AI should take. For each: is there an API? Who approves? What is the blast radius if it is wrong?

3. Decompose one initiative

Pick your highest-priority AI initiative. Walk it down the ladder: Vision - Outcome - Projects - Next Actions.

Let's Continue the Conversation



LinkedIn



Slides + Worksheet

Thank You

Kanishk Mittal

AI Council • May 2026