WEEKLY AI PULSE · VOL. 01

WEEK OF 8 TO 14 JULY 2026

The capability curve keeps steepening. The deployment curve does not.

Six signals from the week. Three model launches. One number worth sitting with. One pattern that changes what boards should be asking on Monday.

IN THIS EDITION

02 One number worth sitting with
03 The story that matters
04 Three things to know
05 The pattern I see

80%
APPS EMBED AN AGENT

88%
PILOTS NEVER SCALE

171%
ROI WHEN THEY DO

ONE NUMBER

Worth sitting with.

88%

of enterprise AI agent pilots never reach production.

AND THE 12% THAT DO

171% average ROI at 12 months.

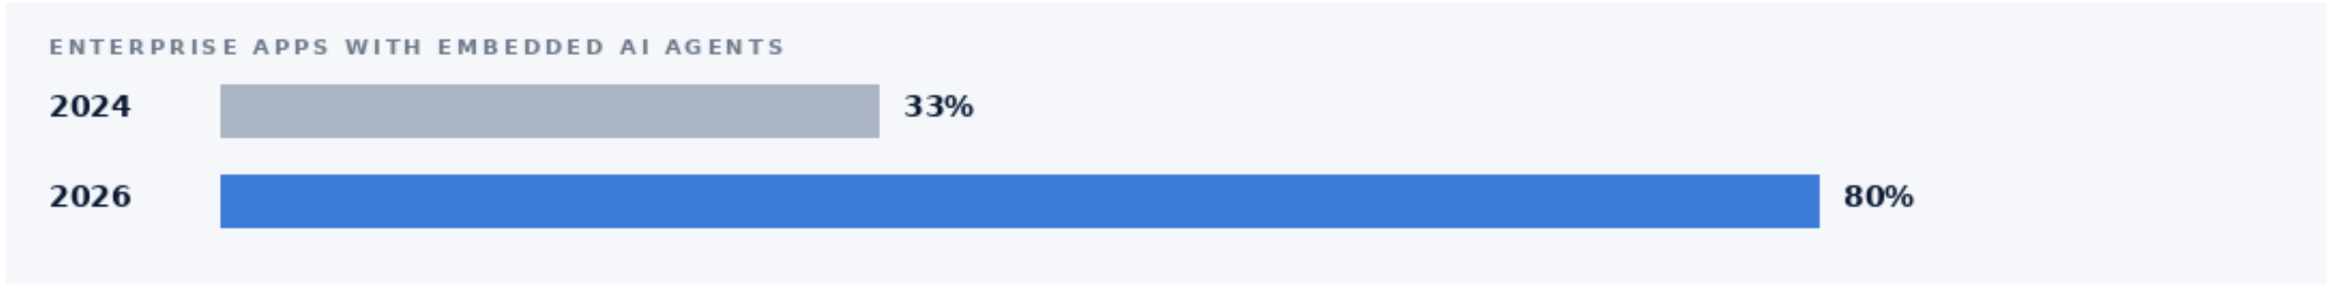
SIGNALS TO WATCH

- → Named owner in the operating model, not just a project sponsor.
- → Line item in the P&L, not a slide in the innovation deck.
- → Governance framework specific to the paradigm (predictive, generative, agentic).

THE STORY THAT MATTERS

AI agents moved to production this year. Governance did not.

80% of production applications now embed at least one AI agent. Up from 33% two years ago. Fastest software category shift in a decade.



95%

OF IT LEADERS

report integration issues at scale.

31%

AT LEAST ONE LIVE

of enterprises have an agent in production.

THE GAP

Capability shipped. Governance stayed in the pilot phase.

THREE THINGS TO KNOW

The signal underneath the noise.

MODEL RELEASES THIS MONTH

Claude Fable 5 | Claude Sonnet 5 | GPT-5.6 | Grok 4.5 | Gemini 3.5 Flash

01 Frontier models got cheaper. Again.

Price-performance for top-tier models has roughly tripled in twelve months. Cost of accessing intelligence is falling fast.

The cost of accessing intelligence keeps falling. The cost of applying it has not moved.

02 Chinese models are enterprise infrastructure.

30 to 46% of enterprise API token usage on US developer platforms now runs on Chinese models. Z.ai's GLM-5.2 grew ~27x in tokens and ~80x in customers in its first week.

Model provenance is now a board conversation, not a procurement footnote.

03 MENA regulation is ahead of Western commentary.

MENA is emerging as one of the world's leading regions for AI governance and sovereign AI. Governments across the region are moving from AI strategies to concrete regulatory frameworks, positioning the Middle East ahead of many markets in operational AI readiness.

The window to build defensible AI programs here is open. It will not stay open long.

THE PATTERN I SEE

Two curves.
Opposite directions.

Every week brings a new frontier model. Cheaper, faster, more capable. That curve is steepening.

Every week also brings evidence that only ~12% of enterprise AI initiatives reach production. That curve is flat.

The strategic question of the next 18 months is not which model to buy. It is which operating discipline to build around it.

Organisations that solve the deployment layer will run three years ahead of the ones still buying licences and calling it a strategy.

THE OPERATOR TEST

Name the paradigm. Name the owner.
Name the workflow. Name the P&L line.

IF YOU RUN AN ORGANISATION

- → Audit which AI paradigms you are actually deploying.
- → Give each paradigm its own owner, KPI and governance.
- → Do not confuse licence utilisation with operating impact.

SEE YOU NEXT MONDAY

Weekly signal. Filtered for operators.

One carousel every Monday. The signals that matter to leaders running organisations, cut from the noise of the week.

COMING UP IN VOL. 02

- → The one metric that separates AI investors from AI collectors
- → How three MENA banks quietly moved past the pilot phase
- → The compliance shift that will hit every European operator

FOLLOW RAVINEL.AI FOR VOL. 02