

INDEPENDENT MARKET ANALYSIS · MID-2026

The No-Code Agent Builder Market

A two-axis independent scoring of 27 enterprise vendors — who leads, who disrupts, and why the no-code persona decides the standings.

Market Shapers


Pace Setters

Pioneers

Specialists

Jukka Niiranen | Power Platform Advisor & Owner

June 2026 · Source: Claude Fable 5 Research · vibes.jukkan.com/ncab-quadrant.html



How to read this report

Independent, not a reprint

Original scoring calibrated in June 2026. It does not reproduce Gartner's or Forrester's placements — it is a transparent, evidence-graded point of view.

Scored for one persona

Every score targets the no-code / business-user buyer. Developer-first strength (AWS AgentCore, Vertex ADK, SAP's pro-code tooling) is deliberately discounted.

Evidence is graded, not assumed

Headline adoption numbers are flagged by source quality. Earnings filings outrank press releases, which outrank marketing claims.

EVIDENCE GRADING

Earnings-grade

SEC filings & earnings calls. Microsoft org counts, Salesforce Agentforce ARR / deal counts.

Highest trust

Disclosed / investor

Company press releases, funding rounds, investor letters. Glean ARR, challenger valuations.

Use with care

Vendor marketing

“Tens of thousands of customers,” “millions of agents,” MAU counts — unaudited and directional only.

Treat as claim

Incumbents who pair distribution with consumption pricing lead the market

Microsoft and Salesforce sit highest on execution; with Google and ServiceNow they form the “Market Shapers” who set the agenda. The most disruptive innovation comes from AI-native challengers — but they lack the install base, governance, and channel to match on execution.

THE BEST-SOURCED STORY IN THE CATEGORY

Two leaders are quantitatively validated by earnings disclosures — everyone else is reacting to them.

01

Distribution beats novelty

Agents win where employees already work — Teams, Slack, the CRM, the ServiceNow portal — not where they are smartest.

02

Governance is table stakes

The entire 2026 product cycle is about agent inventory, permissions, evaluation and cost control — “shadow agents” is the top CIO fear.

03

Pricing is going consumption

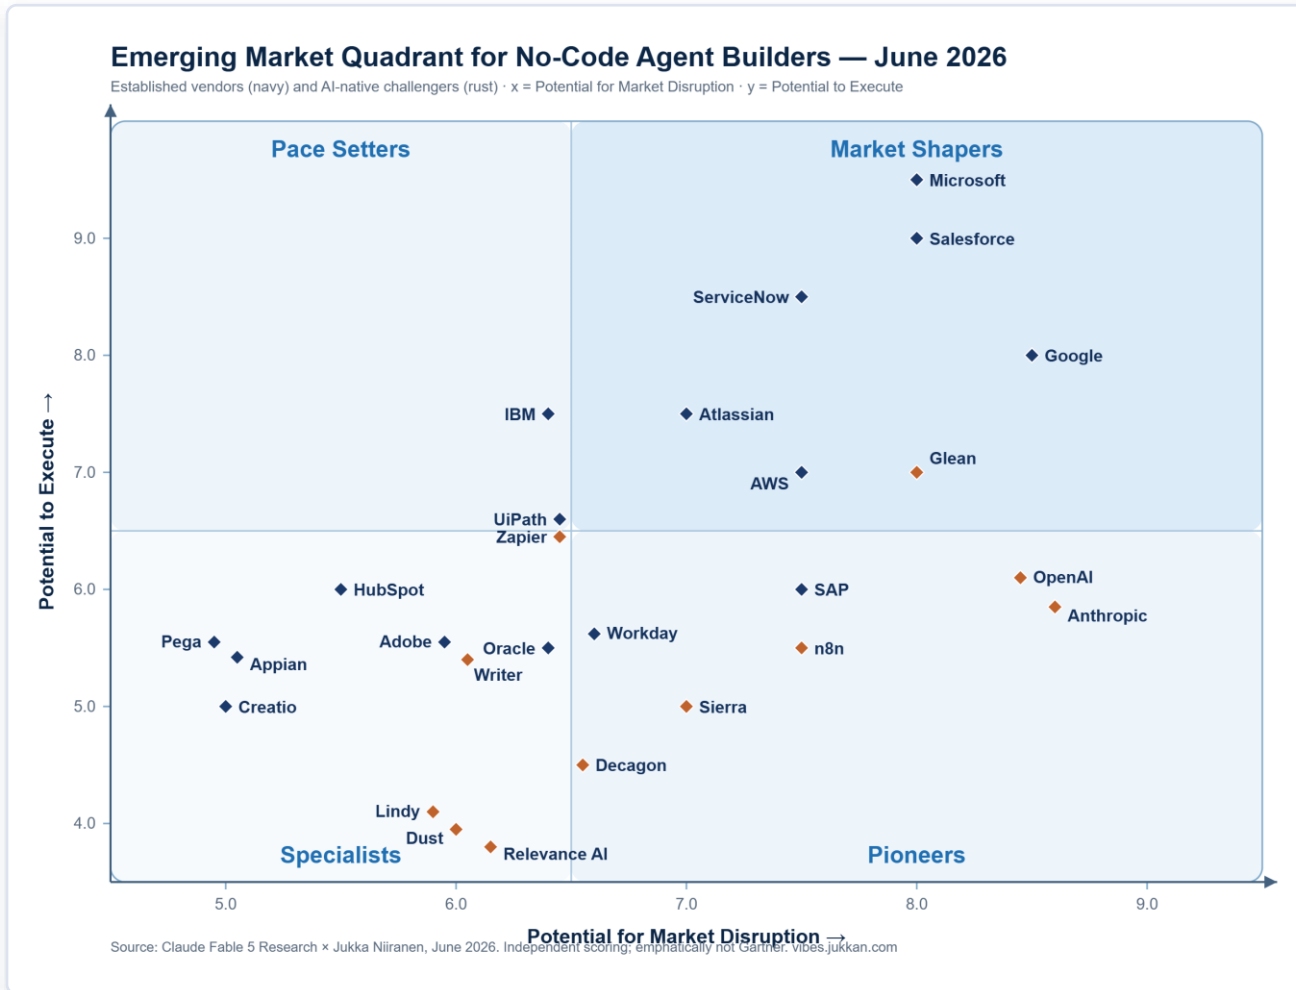
Per-seat licensing is being displaced by per-action / per-credit / per-outcome models — flexible, but with real bill-shock risk.

04

Persona is decisive

Scoring the business-user buyer rewards Microsoft, Salesforce, ServiceNow & Atlassian and discounts developer-first platforms.

Twenty-seven vendors, two axes, four quadrants



7

Market Shapers

High execute · high disrupt. Microsoft, Salesforce, Google, ServiceNow set the pace; Atlassian, AWS & Glean ride the border.

4

Pace Setters

High execute · lower disrupt. IBM, SAP, UiPath, Oracle — strong delivery, repositioning rather than redefining.

10

Pioneers

Lower execute · high disrupt. OpenAI, Anthropic, n8n, Sierra — the innovation edge, gated by install base & governance.

6

Specialists

Lower execute · lower disrupt. HubSpot, Adobe, Pega, Appian — deep in a niche, narrow on general business-user agents.

Axis convention: x = Potential for Market Disruption · y = Potential to Execute. Threshold ≈ 6.5 on each axis.

What each axis measures

Y • Potential to Execute


Can a business user ship a governed agent today?

- Enterprise install base & distribution
- Governance / security / admin depth
- Pricing clarity & connector ecosystem
- Partner network & financial resources
- Evidence of business-user agents in production

X • Potential for Market Disruption

How fast is it redefining how agents get built?


- Pace of innovation & novel capabilities
- Autonomous / multi-agent orchestration, computer use
- MCP / A2A interoperability
- Pricing-model disruption
- Developer & community momentum, mindshare



Threshold ≈ 6.5 / 10 on each axis. Scores are calibrated to the no-code business-user persona — pure pro-code strength does not transfer.

The four who set the agenda

Highest execution paired with high disruption. These four combine the largest install bases with the most aggressive consumption pricing — everyone else reacts to them.



Borderline Shapers: Atlassian (Rovo), AWS (Bedrock AgentCore / Quick Suite) and Glean ride the threshold — strong, but each gated on the strict no-code lens.

Microsoft — the default the others react to

Copilot Studio · Agent 365 · Power Platform

230,000+

organizations on Copilot Studio

~90%

of the Fortune 500 with active low-code agents

\$0.01

per Copilot Credit · \$200/mo = 25,000 credits

\$15 / \$99

Agent 365 per user · E7 “Frontier” suite

Why it leads

- Deepest governance estate — Agent 365 reached GA on May 1, 2026 (\$15/user/mo) with identity, inventory and cost controls.
- Consumption pricing by default — moved from per-message to Copilot Credits; credit consumption up nearly 2x quarter-over-quarter.
- Computer-use agents reached GA in May 2026, extending reach beyond connectors.
- Disruption is capped at 8.0 only because Microsoft is the incumbent others position against.

● Copilot Studio — Agent Builder

“Describe what this agent should do...”

Trigger

Knowledge

Action

● Governed by the admin control plane — identity, permissions, cost & audit

Salesforce — the best-sourced growth story

Agentforce · Agent Builder · Agentforce Command Center

\$1.2B

Agentforce ARR (Q1 FY27 8-K)

+205%

year-over-year ARR growth

29,000+

Agentforce deals closed

3.8B

Agentic Work Units, +111% Q/Q

Why it leads

- A CRM data moat plus Slack distribution puts agents inside the system of record sellers and service reps already use.
- Aggressive Flex Credits consumption (\$500 / 100,000 credits; ~\$0.10 per action) plus a \$2/conversation option.
- ARR and deal counts come from an 8-K filing — the most rigorously sourced adoption story alongside Microsoft.
- Friction: pricing complexity and a Data Cloud dependency raise the cost of entry.

● Agentforce — Agent Builder

“Describe what this agent should do...”



● Governed by the admin control plane — identity, permissions, cost & audit

Google & ServiceNow — the other two shapers

Google

8.0 / 8.5

Gemini Enterprise Agent Platform · Agent Designer

■ Biggest platform move of 2026

Rebranded Vertex AI to the Gemini Enterprise Agent Platform at Cloud Next '26 (Apr 22).

■ No-code for business users

Agent Designer ships inside the Gemini Enterprise app; 200+ models including Anthropic's Claude.

■ Open & funded

Production A2A protocol, Workspace distribution, and a \$750M partner fund.

■ The caveat

Frontier-model pace tops disruption; business-user execution is younger than Microsoft's.

ServiceNow

8.5 / 7.5

Now Assist · AI Agent Studio · Otto

■ The control-tower narrative

Now Assist tracking toward \$1.5B 2026 ACV, up from a prior \$1B target.

■ A conversational front door

Closed the \$2.85B Moveworks acquisition (Dec 15, 2025); rebranded the combined stack as Otto.

■ Governed by design

AI Control Tower governs the estate; natural-language building in AI Agent Studio.

■ The caveat

Workflow-centric and extremely sticky in IT/HR/CS — but bounded to that gravity.

Three forces shaping the market



01

Distribution beats novelty

Incumbents win not because their agents are smartest, but because agents surface where employees already work — Teams, Outlook, Slack, the CRM, the ServiceNow portal. Internal-software adoption is a distribution problem masquerading as a training problem.



02

Governance is table stakes

“Agent sprawl” and “shadow agents” are the dominant CIO anxiety. The entire 2026 product cycle — Agent 365, AI Control Tower, Glean Protect, watsonx control plane — is about inventory, permissions, evaluation and cost forecasting.



03

Pricing converges on consumption

Per-seat licensing is giving way to per-action / credit / conversation / outcome models — flexibility plus “bill shock” risk, which itself spawned new tooling (Microsoft's usage estimator, Salesforce's Digital Wallet).

Per-seat is being structurally displaced



Vendor	Consumption model
Microsoft	Copilot Credits — \$0.01 / credit; \$200/mo = 25,000 credits
Salesforce	Flex Credits — \$500 / 100,000; ~\$0.10 / action; or \$2 / conversation
ServiceNow	Consumption “Assist packs”
Google · AWS · IBM	Usage- / MAU-based metering

THE SINGLE BIGGEST COST LEVER

\$0.01 PAYG per credit

vs.

\$200/mo capacity pack

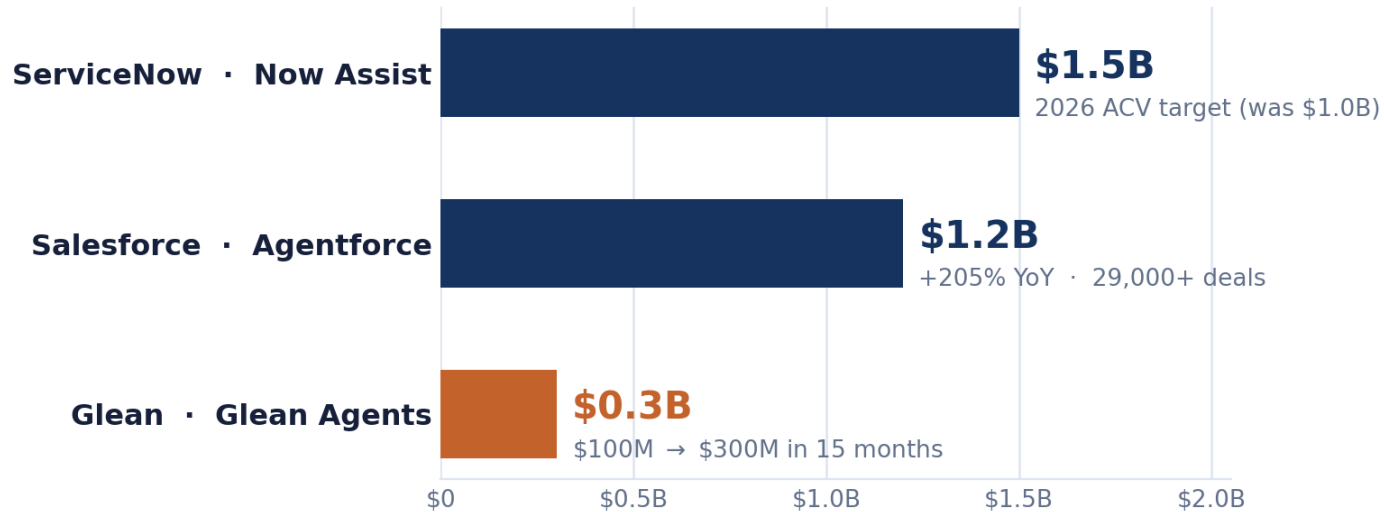
Real-world adoption routinely overshoots forecasts.

Budget for 1.5–3x projected volume, then pre-commit capacity once you have 30 days of actual usage.

The dynamic a Power Platform licensing expert will scrutinize most.

Read the numbers by the quality of their source

Disclosed AI-agent revenue signals



Metrics differ (ARR vs ACV target) and are vendor-disclosed; Salesforce/ServiceNow are earnings-grade, Glean is a company press release.

What is rock-solid

- Microsoft's org count & Fortune 500 penetration (FY26 Q3 earnings).
- Salesforce's Agentforce ARR & deal counts (Q1 FY27 / Q4 FY26 8-Ks).

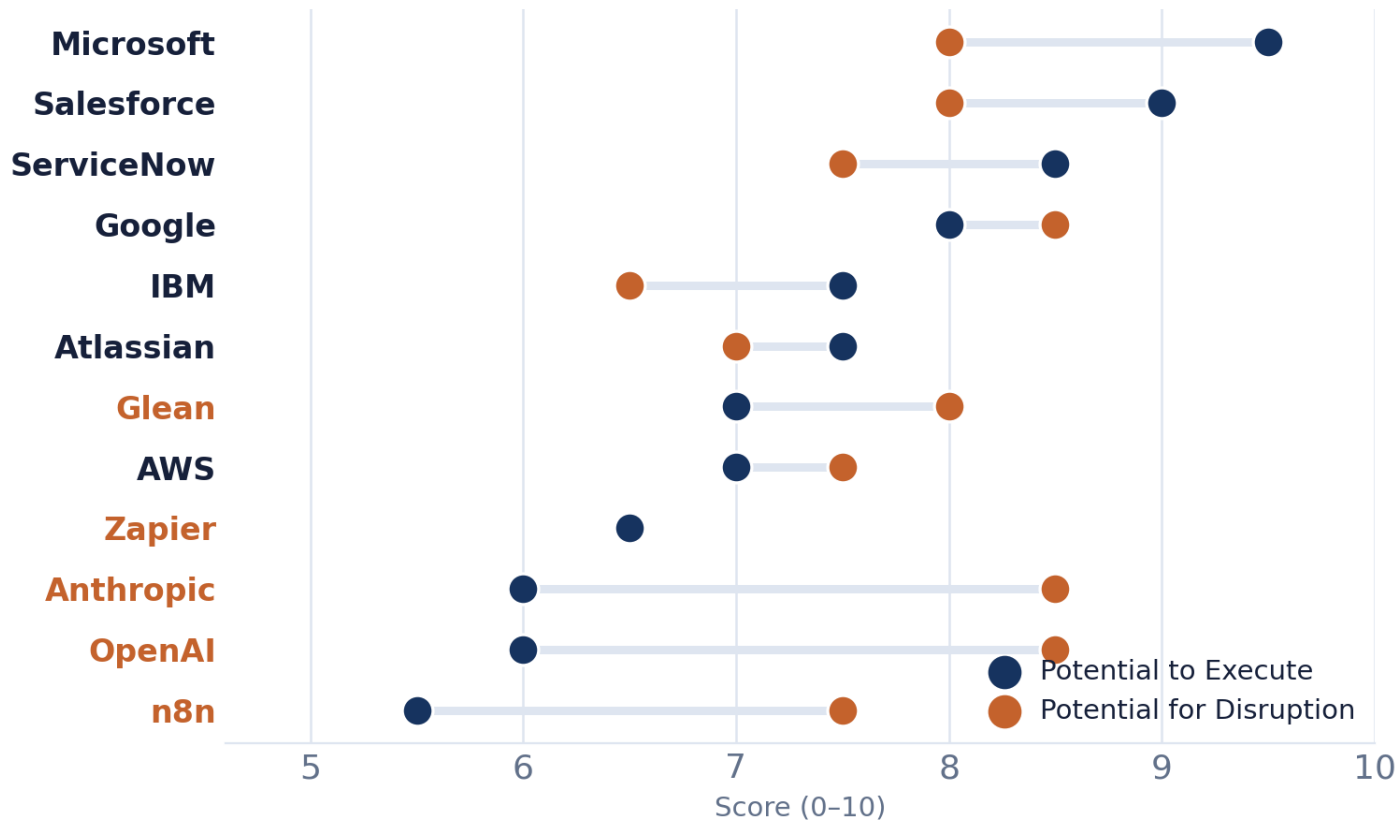
What needs a pinch of salt

- Challenger ARR and valuations are company- or investor-sourced.
- “Tens of thousands of customers,” “millions of agents” and MAU counts are unaudited marketing.
- The circulating “\$750M Q1 Now Assist ACV” figure could not be confirmed in official disclosures.

Glean reached \$300M ARR just 15 months after \$100M — fast, but a press-release figure.

Execution vs. disruption — the defining tension

Execution vs. Disruption — where the gap lives



The pattern

Incumbents cluster top-right on execution; AI-native challengers stretch right on disruption but sink on execution. The wider the bar, the bigger the gap between what a vendor can imagine and what a business user can safely ship today.

■ Microsoft & Salesforce

Tight bars — execution and disruption move together at the top.

■ OpenAI & Anthropic

The widest gap: disruption 8.5, execution only 6.0 — thin enterprise governance.


■ n8n

Disruptive on price & control, but developer-leaning keeps no-code execution low.

Incumbents (navy labels) lead on execution; AI-native challengers (rust labels) lead on disruption but trail on execution.

Scaling fast — but narrow

The most disruptive innovation comes from AI-native challengers and workflow players. Most lack the install base, governance maturity and channel to score high on execution.




Glean — \$300M ARR, natural-language Auto Mode builder, strong Glean Protect governance; lacks a captive install base.

n8n — \$5.2B valuation (SAP-backed), 230k+ active users, multi-agent orchestration; more developer-leaning than true no-code.

Sierra (~\$15.8B) & Decagon (~\$4.5B) — genuine production CX deployments, but guided/configured rather than open citizen-developer builders.

⚠️ Churn signal: OpenAI's Agent Builder is being wound down (sunset Nov 30, 2026) in favor of Workspace Agents + the Agents SDK.



Private valuations from funding rounds / investor disclosures (2025-26). Lovable is an adjacent app-builder, not a pure NCAB vendor.

Most contested placements

Google

Shaper vs Pioneer

Frontier disruption is real, but business-user execution is younger. Disclosed seat/agent counts would lock in Shaper.

AWS

Execution score

AgentCore is best-in-class infrastructure but developer-first; Quick Suite is unproven for citizen developers.

OpenAI / Anthropic

Execution score

Huge mindshare, thin enterprise governance. The Agent Builder sunset injects real uncertainty about no-code commitment.

Glean

Shaper vs Pace Setter

\$300M ARR argues up; no captive install base argues for caution. Fortune 500 penetration is the tiebreaker.

n8n vs Zapier

No-code purity

n8n scores higher on disruption; Zapier is the truer business-user tool. Reasonable analysts could swap their x-values.

SAP Joule Studio

Persona mismatch

Agent Builder is GA but developer-skewed. The managed, zero-infra no-code version only GAs in Q3 2026.

What to do about it

1 For a Microsoft-centric enterprise

Standardize on Copilot Studio + Agent 365. Model costs in Copilot Credits from day one — the \$0.01 PAYG vs \$200/mo pack decision is the biggest lever. Track the E7 “Frontier” suite (\$99/user/mo) as the consolidation SKU.

2 For Salesforce / ServiceNow shops

Use the native builder — data gravity and unified governance outweigh best-of-breed. Negotiate Flex Credits / Assist-pack commits by modeling actions-per-workflow; \$2/conversation only wins above ~20 actions per interaction.

3 For horizontal needs

Evaluate Glean (knowledge layer) and n8n or Zapier (orchestration) as complements, not replacements. Pick n8n where you own engineering; pick Zapier where non-technical operators own the workflow.

4 Treat OpenAI & Anthropic as suppliers

Pilot Cework / AgentKit for prosumer productivity, but don't anchor enterprise governance on them until audit features mature — and factor in OpenAI's Agent Builder sunset.

Benchmarks that change the strategy: audited business-user agent counts in production · consumption pricing stabilizing into per-outcome rates · MCP/A2A maturing enough to govern a multi-vendor estate from one control plane.

APPENDIX

The full vendor scoreboard

All 27 scored vendors. *y* = Potential to Execute · *x* = Potential for Market Disruption.

Vendor	y	x	Quadrant
Microsoft	9.5	8.0	Shaper
Salesforce	9.0	8.0	Shaper
ServiceNow	8.5	7.5	Shaper
Google	8.0	8.5	Shaper
Atlassian	7.5	7.0	Shaper†
AWS	7.0	7.5	Shaper†
Glean	7.0	8.0	Shaper†
IBM	7.5	6.5	Pace Setter
UiPath	6.5	7.0	Pace Setter
Oracle	5.5	6.5	Pace Setter
SAP	6.0	7.5	Pace Setter
OpenAI	6.0	8.5	Pioneer
Anthropic	6.0	8.5	Pioneer
n8n	5.5	7.5	Pioneer

Vendor	y	x	Quadrant
Zapier	6.5	6.5	Pioneer†
Sierra	5.0	7.0	Pioneer
Decagon	4.5	6.5	Pioneer
Writer	5.5	6.0	Pioneer†
Workday	5.5	6.5	Specialist†
Lindy	4.0	6.0	Pioneer
Dust	4.0	6.0	Pioneer
Relevance AI	4.0	6.0	Pioneer
HubSpot	6.0	5.5	Specialist
Adobe	5.5	6.0	Specialist†
Pega	5.5	5.0	Specialist
Appian	5.5	5.0	Specialist
Creatio	5.0	5.0	Specialist

† borderline placement. Source: Claude Fable 5 Research · Jukka Niiranen, June 2026.

Caveats & how to use this

Vendor-marketing risk

Many figures are company-reported and unaudited. The most reliable are Microsoft's earnings figures and Salesforce's 8-K ARR/deal counts; challenger ARR and valuations are company- or investor-sourced.

Short half-life

The category moves monthly — OpenAI's Agent Builder sunset, SAP's Q3 managed GA, ServiceNow's Otto rebrand. Placements should be revisited each quarter.

Entity status

Moveworks is scored as absorbed into ServiceNow (closed Dec 15, 2025). Lovable is adjacent (an app-builder) — its strategic cloud/model partner is Google Cloud, not Microsoft.

Scores are judgment calls

Calibrated to the no-code business-user persona as of June 2026. They are not Gartner's and do not reproduce its anonymized placements.

Scope discipline

Pure pro-code frameworks (LangChain/LangGraph, Agent Framework/AutoGen, CrewAI, Strands, Vertex ADK) are excluded except as context.

“The category moves monthly.”

Distribution, governance and consumption pricing are the three axes that will decide the next two quarters. Re-score when a vendor discloses audited business-user agent counts — that is the signal that turns best-of-breed Pioneers into safe bets alongside the incumbent.

Jukka Niiranen · Power Platform Advisor & Owner

Independent analysis · June 2026 · Source: Claude Fable 5 Research

vibes.jukkan.com/ncab-quadrant.html

