

# The Rent and the Reach

---

Power Platform against AI-assisted development, and why the low-code bargain stopped clearing

---

**David Soden**

Independent analyst and practitioner, enterprise automation and applied AI  
[davidsoden.com/market-assessment](https://davidsoden.com/market-assessment)

## About this report

Every substantive claim in this report carries an evidence classification and, where applicable, a numbered source. Facts are separated from vendor claims, third-party estimates, the author’s assessments, and untested hypotheses. Forward-looking sections use scenarios with observable tripwires rather than point forecasts. Stated assumptions are listed before the analysis begins.

**PUBLISHED FOR PUBLIC READING | SHARE FREELY, UNMODIFIED, WITH ATTRIBUTION**

© 2026 David Soden, [davidsoden.com/market-assessment](https://davidsoden.com/market-assessment). All rights reserved. You may share this file complete and unmodified, and quote short passages with attribution. Republishing under another name, excerpting into a commercial deliverable, resale, and use as machine learning training data require written consent. Full terms appear under Rights and Permitted Use on the final page.

Market Assessment reports are published and commissioned at [davidsoden.com/market-assessment](https://davidsoden.com/market-assessment). This document is analysis, not investment, legal, or procurement advice.

## Scope and evidentiary standard

This report assesses whether Microsoft Power Platform still earns its price for greenfield internal application development, measured against AI-assisted development performed by non-developers. It covers licensing as published, the range of artifacts each path can produce, the published evidence on AI-assisted code quality and speed, and the governance question that has historically justified the platform premium.

Two things about this category make published numbers unusually weak. First, almost every quantified productivity claim on the low-code side traces to research the platform vendor commissioned, and almost every quantified claim on the AI side comes from firms selling code-quality or security tooling. Second, and more seriously, nobody has run a controlled comparison of the two paths on the same task. The central question in this market has never been measured. That absence is a finding, not a gap to paper over.

Every claim below carries a label. Where the evidence is a vendor’s own account, it says so. Where the judgment is the author’s, it says that too, and the reasoning is shown so it can be argued with.

| Label        | Definition                                                                                                               | How to treat it                                                                                                 |
|--------------|--------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| FACT         | Verifiable in the cited primary source: a filing, press release, official documentation, or standards body announcement. | Reliable as stated. Check the source if the exact wording matters.                                              |
| VENDOR CLAIM | Reported by a vendor about its own business or product. Self-measured and not audited by any third party.                | The vendor said it. Directionally useful, not a measurement. Definitions of the metric are usually undisclosed. |

| Label                | Definition                                                                                                                         | How to treat it                                                                                                          |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| THIRD-PARTY ESTIMATE | A research firm forecast or survey result. Methodology is proprietary and generally not reproducible.                              | Use for direction and order of magnitude only. Do not build a model on it.                                               |
| ASSESSMENT           | The author's reasoned judgment from the cited evidence. Falsifiable, and the reasoning is shown so it can be argued with.          | This is analysis, not reporting. Disagree with it where your evidence differs.                                           |
| HYPOTHESIS           | A proposition offered for testing. Not currently supported by sufficient evidence to assert.                                       | Treat as a research question, not a conclusion. Each one names what would confirm or refute it.                          |
| SCENARIO             | A structured future state with a stated subjective probability. The probability is the author's judgment, not a calculated figure. | Use the leading indicators attached to each, not the probability. The indicators are observable; the probability is not. |

## Assumptions this analysis rests on

Three premises are assumed rather than shown. Each names what fails if it is wrong.

First, that published list pricing approximates what mid-market buyers actually pay. Large enterprises negotiate, and a heavily discounted Enterprise Agreement changes the arithmetic in section six. If this is wrong, the cost comparison narrows but the addressable-surface argument in section five is untouched.

Second, that current AI coding agents remain available at roughly current capability and price through the assessment window. If agent pricing rises sharply or capability plateaus below present levels, the speed and reach case weakens and the platform's managed-service premium looks better.

Third, that the author's shipped portfolio is representative of a meaningful share of internal tooling demand, rather than an unusual set of edge cases. This is the weakest of the three and it is examined directly rather than assumed away in section five. If it is wrong, the addressable-surface finding shrinks to a niche observation.

## The call

**ASSESSMENT** Power Platform's original bargain, faster delivery by non-developers in exchange for per-seat rent, no longer clears for greenfield internal tooling. The rent rose while the addressable surface stayed fixed, and AI-assisted development now reaches artifacts the platform cannot produce at any price. I would change position if Microsoft decoupled Power Apps pricing from per-user seats, or if a controlled study showed non-developers shipping production internal apps faster in a canvas designer than with an agent. Where Dataverse is the system of

record, or where an organization genuinely cannot trust its builders, the platform premium still survives. That carve-out is real and it is not small.

*My judgment on the evidence assembled here, nothing more. There is data I can't see and data I may have missed. New evidence can move me, including to the opposite position, and I reserve the right to move.*

## Executive summary

Low-code was sold on a specific trade: business technologists ship working applications without waiting for scarce, expensive professional developers, and in exchange the organization pays per-seat rent to the platform. That trade was defensible in 2016. Four things have changed since, and three of them moved against the platform.

### FINDING 1 The entry price roughly doubled while the product did not.

**FACT** Microsoft's public Power Apps pricing page now lists Premium at \$20.00 per user per month and a 2,000-seat rate of \$12.00. The \$5 per app plan that was the cheap way in no longer appears on it, and the metered alternative runs at roughly twice the per app rate. <sup>[2] [4] [5]</sup>

### FINDING 2 Microsoft is withdrawing an entitlement mid-flight, not just raising a price.

**FACT** Seeded AI Builder credits inside Power Apps Premium, Power Automate Premium and Dynamics 365 licenses are removed on 1 November 2026 for all new and existing customers, including Enterprise Agreement holders. There is no conversion to Copilot Credits. Features already running in production stop unless new capacity is purchased. <sup>[1]</sup>

### FINDING 3 The strongest argument is reach, not speed.

**ASSESSMENT** Of nine products the author shipped as a non-developer, seven cannot be produced by Power Platform at any price and two require significant workaround. Signed desktop applications, self-hosted container services and local-only data processing are outside the platform's addressable surface entirely. Price is a second-order argument when the first-order answer is that the tool cannot make the thing. <sup>[15]</sup>

### FINDING 4 The productivity question this entire market is priced on has never been measured.

**ASSESSMENT** No controlled study compares building the same internal application in a canvas designer against building it with a coding agent. The quantified claims in circulation come from research the platform vendor commissioned. When neither side has a measurement, the burden falls on the party charging rent. <sup>[6] [7]</sup>

### FINDING 5 AI-assisted code has a real quality problem, and it is not the problem the platform solves.

**THIRD-PARTY ESTIMATE** Independent and vendor research finds refactoring in sharp decline and security defects in a large minority of generated code. This is a genuine liability. It is also a discipline problem addressed by process harnesses and shared skills, not something a per-seat license fixes. <sup>[8] [9] [10]</sup>

### FINDING 6 Guardrails are no longer proprietary to the platform.

**FACT** The Agent Skills format was published as an open standard in December 2025 and donated to a Linux Foundation-backed body, with Microsoft, GitHub, OpenAI and others adopting it. Structured development harnesses now carry tens of thousands of GitHub stars. Encoded process is available outside a licensed platform in a way it was not two years ago. <sup>[11] [16]</sup>

## The bargain the platform was sold on

**ASSESSMENT** The 2016 case for low-code had nothing to do with code quality and everything to do with queue position. Professional developers were scarce and expensive, internal tooling sat at the bottom of their backlog, and a business technologist who understood the process could not get a simple approval app built inside a fiscal year. Low-code sold the ability to skip that queue. The per-seat cost was framed as cheap against the alternative, which was waiting.

**ASSESSMENT** That framing had a hidden assumption inside it: that the alternative to the platform was a professional developer. Everything in the pricing model follows from that comparison. Twenty dollars a user per month is trivial against a contract developer's day rate, and Microsoft has priced accordingly for a decade.

**ASSESSMENT** The assumption is what broke. The alternative to the platform for a business technologist is no longer a professional developer. It is the same business technologist, working with an agent, producing a wider range of artifacts. The comparison the price was set against has been replaced, and the price has not moved to reflect it.

**VENDOR CLAIM** Microsoft's own framing acknowledges the shift. At the Power Platform Community Conference in October 2025 the company's leadership told the audience that low-code as it had been known was dead and being reborn with AI at its core, alongside a claim that Copilot Studio usage grew from 50,000 to 200,000 organizations in a year. <sup>[13]</sup>

**ASSESSMENT** Read that growth figure carefully. It counts organizations with any usage, from an undisclosed base, on a product bundled into existing licenses. It measures distribution, not adoption, and certainly not retention. A vendor reporting organization counts rather than seats, revenue, or workloads in production is choosing the most flattering available denominator. <sup>[13]</sup>

## What the licensing actually did in 2026


**FACT** The Power Apps per app plan, historically \$5 per user per app per month, reached end of sale on the standard price list on 2 January 2026. Practitioner licensing specialists report that Enterprise Agreement customers may continue to renew, that the SKU returned to Cloud Solution Provider price lists in the spring, and that MPSA customers lost it outright. <sup>[4] [5]</sup>

**FACT** Microsoft's current public Power Apps pricing page lists Power Apps Premium at \$20.00 per user per month paid yearly, a \$12.00 per user per month rate at a 2,000-seat minimum, and Copilot Studio at \$200.00 for 25,000 Copilot Credits per month. No per app plan appears on it. <sup>[2]</sup>

**ASSESSMENT** The practical effect for a new mid-market buyer is that the cheapest supported way to put one application in front of one hundred occasional users roughly doubled, and the

fallback is a metered rate that charges by active user per app rather than by seat. Metering is more honest for genuinely occasional use. It is also harder to budget, which is a real cost that never appears on a price list. <sup>[14]</sup>

**FIGURE 1** Entry cost per user, per month, by licensing path



**FACT** Only the two Premium figures appear on Microsoft's current public pricing page. The per app and metered rates are reported by licensing practitioners, not published by Microsoft in the same place, and regional pricing varies. <sup>[2] [4] [5] [14]</sup>

**FACT** Separately, Microsoft is removing seeded AI Builder credits from Power Apps Premium, Power Automate Premium and Dynamics 365 licenses on 1 November 2026. The change applies to all new and existing customers including those on Enterprise Agreements. There is no automatic conversion of AI Builder credits into Copilot Credits, and Microsoft's documentation states plainly that after that date the seeded credits are unavailable and there is no transition. <sup>[1]</sup>

**ASSESSMENT** This is materially different from a price rise, and buyers should treat it differently. A price rise is visible at renewal and can be negotiated. An entitlement withdrawal converts a capability that was inside a license the customer already bought into a separate consumption line item, and it does so on Microsoft's calendar rather than the customer's. Contracts that started before 1 November 2025 are honored for their term, which means the exposure lands unevenly across a customer base and is easy to miss until a production flow stops. <sup>[1]</sup>

**ASSESSMENT** For an investor, the read is that Microsoft is converting Power Platform from a seat-licensed product into a consumption-metered one. That transition usually raises revenue per customer and it also raises churn risk, because consumption pricing makes the bill visible monthly rather than annually. Buyers who never audited a seat count will audit a credit burn.

## The addressable surface

**ASSESSMENT** The speed argument is contestable. The reach argument is close to settled, and it is the one that should decide most buying conversations. Power Platform produces a specific and bounded set of artifacts: canvas and model-driven applications inside the Microsoft tenant, automation flows, Power Pages sites, and Copilot Studio agents. It does not produce signed desktop applications, self-hosted container services, browser extensions, or software that processes data entirely on a local machine.


**ASSESSMENT** That boundary was largely invisible while the alternative was a professional development queue, because the same constraint applied on both sides for a business technologist. It becomes highly visible the moment the same person can produce those artifacts directly. <sup>[15]</sup>

**FACT** The author's published portfolio contains nine shipped products: four code-signed desktop applications for Mac and Windows, three self-hosted Docker services, a browser extension, and model-backed web applications. Several process data entirely locally by design, including clinical health data on the user's own machine. <sup>[15]</sup>

**FACT** The nearest platform equivalent for external-facing work, Power Pages, is licensed at \$200 per month for each pack of 100 authenticated users per site, assigned per environment. Users needing access to a second site consume a second pack. <sup>[3]</sup>

**ASSESSMENT** That pricing shape matters more than the headline number. It scales with audience size rather than with value delivered, which means a successful portal is punished for succeeding. A self-hosted equivalent carries the same infrastructure cost at one hundred users as at ten thousand. <sup>[3]</sup>

| Artifact class                               | Power Platform equivalent                                             | Assessment                                                                                   |
|----------------------------------------------|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Code-signed desktop app, Mac and Windows     | None. The platform has no desktop distribution target.                | Outside the addressable surface. No price would change this.                                 |
| Self-hosted Docker service                   | None. Power Platform is tenant-hosted SaaS.                           | Outside the addressable surface. Data residency and self-hosting requirements cannot be met. |
| Browser extension                            | None.                                                                 | Outside the addressable surface.                                                             |
| Local-only data processing                   | None. Dataverse is cloud-resident.                                    | Outside the addressable surface, and the constraint is architectural rather than commercial. |
| Externally facing portal with authentication | Power Pages, at \$200 per 100 authenticated users per site per month. | Buildable, at a cost that scales with audience rather than with value delivered.             |
| Model-backed internal web app                | Power Apps plus Copilot Studio credits.                               | Buildable with workaround. Credit consumption is the budget risk, not the seat count.        |

**FIGURE 2** Nine shipped products, classified against the Power Platform addressable surface

**ASSESSMENT** The author's classification of his own portfolio, from published product descriptions. A single practitioner's output is not a market sample, and someone building only tenant-resident line-of-business apps would classify very differently. <sup>[15]</sup>

**ASSESSMENT** The honest counter to this finding is selection. A practitioner who builds desktop and self-hosted software will naturally produce a portfolio that low-code cannot match, and that says as much about the practitioner as about the platform. The counter to the counter is that the portfolio was not built to prove a point about licensing, and that the constraint runs one way: the agent path can produce tenant-resident business applications, while the platform path cannot produce signed binaries. Asymmetric capability is not a matter of taste. <sup>[15]</sup>

**HYPOTHESIS** If the boundary matters as much as this section argues, organizations standardized on Power Platform should show a persistent shadow backlog of requests that were refused on platform-capability grounds rather than on priority grounds. Confirming evidence would be a survey of Power Platform centers of excellence reporting the share of intake rejected as out of scope. Refuting evidence would be intake data showing that essentially all internal demand is satisfiable within the platform's artifact set.

## What the productivity evidence actually supports

**ASSESSMENT** The speed claim is where this argument is weakest, and it is weak on both sides. It deserves to be stated precisely rather than asserted.

**FACT** The most cited controlled study on AI-assisted development is METR's randomized trial published in July 2025. Sixteen experienced open-source developers worked across 246 tasks in repositories they had contributed to for roughly five years, averaging over 22,000 GitHub stars and a million lines of code. Developers took 19 percent longer on tasks where AI tools were permitted, while estimating afterwards that the tools had made them 20 percent faster. <sup>[6] [17]</sup>

**ASSESSMENT** That study is frequently quoted as though it settles whether AI-assisted development is faster. It does not, and it was never designed to. It measured expert maintainers doing bug-fix and refactoring work inside large codebases they already knew intimately. That is close to the worst case for an assistant, because the human's context advantage is at its maximum and the task is comprehension-bound rather than production-bound. It says nothing about a non-developer building a new internal application from nothing. <sup>[6]</sup>


**FACT** METR itself has since described the result as historical and has changed its experiment design, reporting that too many participants were unwilling to work without AI tooling to sustain the original protocol. <sup>[7]</sup>

**ASSESSMENT** The corresponding claim on the low-code side is no better supported. The multiples in circulation for low-code delivery speed originate in research commissioned by the platform vendor. A vendor-funded multiple and a practitioner's field estimate are the same class of evidence: an unmeasured claim with an interested party behind it. The difference is that only one of them is being used to justify a recurring per-seat charge.

**ASSESSMENT** The author's own field observation, offered as judgment rather than measurement: across nine shipped products, most were built in under a day and the worst case ran to a couple of days. The same scope on Power Platform would have run to weeks or months where it was possible at all, and on a conventional development path in 2017 to many months. The author is not a professional developer and does not write syntax unaided, which is precisely the population low-code was built and priced for. <sup>[15]</sup>

**HYPOTHESIS** The comparison that would settle this has not been run: one operator of the business-technologist profile, one specification, both paths, wall-clock time to something running in production, repeated across enough tasks to matter. Until someone runs it, every multiple quoted in this market on either side is a practitioner estimate wearing a number.

**THIRD-PARTY ESTIMATE** On quality rather than speed, the picture is less flattering to the AI path. GitClear's longitudinal analysis of code changes reports refactored code falling from roughly 21 percent of changed lines in 2022 to 13 percent in 2023 and 3.8 percent year to date in 2026, alongside block duplication at the highest level in its record. <sup>[8]</sup>

**FIGURE 3** Refactored code as a share of changed lines

**THIRD-PARTY ESTIMATE** GitClear's own analysis of repositories it observes. The firm sells code-quality analytics and has a commercial interest in the finding. The sample is self-selected and the intervening years are not shown here because the published series mixes measures. <sup>[8]</sup>

**THIRD-PARTY ESTIMATE** Veracode's GenAI code security research, covering more than 100 models across 80 curated tasks, reports that when a model could choose between a secure and an insecure implementation it chose the insecure one in roughly 45 percent of cases, with Java worst at 72 percent. Its spring 2026 update reports that security performance has not improved in step with general model capability. <sup>[9] [10]</sup>

**ASSESSMENT** Both findings are real and neither rescues the platform. A per-seat license does not refactor anything, and Power Platform applications have their own well-known maintainability failure mode in undocumented canvas logic that only its author understands. The correct response to generated-code debt is process, review and testing discipline. That is a different purchase from a platform subscription, and it is the subject of the next section.

## Guardrails without a platform

**ASSESSMENT** The most serious remaining argument for the platform is governance. Managed environments, data loss prevention policies, solution layering and tenant-level administration are real capabilities, and an organization that genuinely cannot trust what its builders will do has a reason to pay for them. This is the part of the platform's value that survives the analysis above, and buyers should not talk themselves out of it.

**FACT** What has changed is that encoded process is no longer proprietary. The Agent Skills format was published as an open standard in December 2025 and donated to a Linux Foundation-backed foundation, with Microsoft adopting it across VS Code and GitHub alongside OpenAI, Atlassian, Figma and others. <sup>[11]</sup>

**FACT** Structured development harnesses now carry meaningful adoption. The BMAD Method, an open-source framework of specialized agent personas and quality gates, is at roughly 49,000 GitHub stars and 5,700 forks. Public skill directories report over 1.1 million cumulative installs across 21 supported agents. <sup>[16] [12]</sup>


**ASSESSMENT** The distinction that matters, and it cuts against the enthusiasts: skills and harnesses encode intent, they do not enforce it. A managed environment can stop a maker from connecting a production database to a personal storage account. A skill file cannot, because the operator can ignore it. That is a real difference in kind, not degree, and any argument that treats them as equivalent is wrong.

**ASSESSMENT** The right conclusion is narrower and still damaging to the platform. Enforcement belongs at the infrastructure layer, where it always did: repository permissions, deployment pipelines, secret management and network policy. Those controls exist independently of any application platform and most enterprises already run them for everything else. What the platform uniquely provided was enforcement bundled with authoring, priced per seat. Unbundling those two is now possible, and the bundle was always the expensive part.

**ASSESSMENT** There is a second-order effect worth naming for investors. Shared, versioned skills spread practice faster than any vendor certification program, and they spread bad practice at the same rate. A public directory with an audit function is a supply chain, with the exposure that implies. Treating an installed skill as trusted code is the same mistake as treating an npm package that way, and the industry has already learned that lesson once. <sup>[12]</sup>

## Cost over five years

**ASSESSMENT** The following comparison is deliberately crude and favors the argument this report is making. It counts license and infrastructure cost only. It excludes build labor, maintenance labor, security review and the cost of a failed project, all of which are substantial and none of which are obviously lower on either path. A reader who believes the agent path carries higher maintenance labor can add that column and reach a different answer using nothing but material already in this document.

**FIGURE 4** Five-year license and infrastructure cost, 200 internal users

**ASSESSMENT** Author's calculation from published rates at 200 users over 60 months. Infrastructure figure assumes \$150 per month of compute and storage. Excludes all labor on both sides, which is the largest omitted cost and the one most likely to narrow the gap. <sup>[2]</sup>

**ASSESSMENT** The gap is wide enough that the omitted labor cost would have to be extraordinary to close it, and the direction of the platform's pricing changes is widening it rather than narrowing it. That said, the comparison is least reliable exactly where organizations are largest, because a negotiated Enterprise Agreement can move the platform figures substantially and those terms are not public.

**ASSESSMENT** The one place the arithmetic reverses is Dataverse. Where it is genuinely the system of record rather than a database of convenience, the migration cost, the integration surface and the risk of losing referential integrity across Dynamics 365 make a rebuild a bad trade regardless of licensing. Organizations should separate these two questions rather than treating a Power Platform estate as one decision.

## Scenarios to 2029

**ASSESSMENT** Point forecasts are worthless in a category moving this fast. What follows is three structured futures with subjective probabilities. The probabilities are the least reliable content in this document. The earliest visible signs attached to each are the most useful, because anyone can watch for them without private information.

### **SCENARIO A** Consumption absorbs the seat 50 percent

Microsoft completes the shift from per-user licensing to credit consumption across Power Platform, folding app hosting into the same meter as Copilot. Total spend rises for heavy users and falls for occasional ones. The per-seat argument becomes moot because there are no seats.

**Consequence.** Buyers lose the ability to forecast spend annually, and finance functions start treating internal tooling as a variable cost. Vendors selling seat-based competitors have to follow. The addressable-surface gap remains untouched, so the reach argument in section five survives this scenario intact.

**Earliest visible sign.** Power Apps list pricing on Microsoft's public page is restated in credits rather than per user per month, or a Power Apps consumption meter appears in the Power Platform admin center alongside the existing Copilot capacity tooling.

**SCENARIO B   The platform becomes an agent host   30 percent**

Power Platform stops competing on authoring and repositions as the governed runtime for agents built elsewhere, keeping the managed environment, data loss prevention and tenant administration while conceding the build experience to coding agents.

**Consequence.** This is the outcome that preserves the most value for existing customers, because the governance layer is the part that genuinely does not exist outside the platform. It also concedes the argument this report makes about authoring, and shrinks the licensable surface to administration.

**Earliest visible sign.** Microsoft ships first-class deployment of externally authored applications into managed environments, or Power Platform admin tooling is unbundled and sold separately from Power Apps authoring.

**SCENARIO C   Generated-code debt forces a correction   20 percent**

The maintainability and security findings compound into visible, expensive failures at enough organizations that boards impose restrictions on agent-authored production software, and governed platforms recover on risk grounds rather than on capability.

**Consequence.** The platform premium is re-justified, but as insurance rather than productivity, and priced accordingly. Buyers who rebuilt aggressively in 2026 and 2027 without process discipline carry the cost. This is the scenario in which the author is wrong, and it is not remote.

**Earliest visible sign.** A major regulator or a large enterprise publicly attributes a breach or a material outage to agent-authored code, or a Fortune 500 issues a formal internal prohibition on agent-authored production systems.

Leading indicators

Each of the following is observable without private information. They are more useful than the probabilities above.

| If this happens                                                | It means                                      | Moves toward                                 | Where to watch                                    |
|----------------------------------------------------------------|-----------------------------------------------|----------------------------------------------|---------------------------------------------------|
| Power Apps per app plan is restored to the standard price list | Microsoft saw measurable entry-tier attrition | Weakens the pricing argument in section four | Microsoft public pricing page and CSP price lists |

| If this happens                                                         | It means                                            | Moves toward                               | Where to watch                                 |
|-------------------------------------------------------------------------|-----------------------------------------------------|--------------------------------------------|------------------------------------------------|
| <b>Seeded AI Builder credit removal is delayed past 1 November 2026</b> | Customer pushback landed                            | Weakens the entitlement-withdrawal finding | Microsoft Learn end of AI Builder credits page |
| <b>Power Apps pricing is restated in credits</b>                        | Consumption transition is complete                  | Scenario A                                 | Microsoft public pricing page                  |
| <b>A controlled study of non-developers on both paths is published</b>  | The central question finally gets measured          | Resolves finding 4 in either direction     | METR, academic software engineering venues     |
| <b>A named enterprise attributes a breach to agent-authored code</b>    | Risk narrative hardens                              | Scenario C                                 | Regulatory filings, incident disclosures       |
| <b>Microsoft ships governed deployment of externally authored apps</b>  | Repositioning as runtime rather than authoring tool | Scenario B                                 | Power Platform release plans                   |

## Implications

### For buyers and IT leaders

**ASSESSMENT** Split the estate before deciding anything. Applications where Dataverse is the system of record, or which are deeply wired into Dynamics 365, are a different decision from standalone canvas applications built on SharePoint or SQL. The first group is expensive and risky to move. The second group is frequently a few days of work to rebuild and carries no ongoing seat cost afterwards.

**ASSESSMENT** Then audit the 1 November 2026 exposure specifically, because it is the one change with a hard date attached. Identify every production flow consuming seeded AI Builder credits and determine whether the contract that covers it started before 1 November 2025. Customers on newer contracts lose the entitlement outright.<sup>[1]</sup>

**ASSESSMENT** Do not treat agent-assisted development as governance-free. The infrastructure controls have to exist first: repository permissions, deployment approvals, secret management, dependency review. An organization that cannot run those disciplines today should keep paying for a managed platform, and should be honest with itself about why.

### For investors

**ASSESSMENT** The metric to watch on Power Platform is not organization counts, which Microsoft prefers to report, but revenue per customer and the mix between seats and consumption. The transition to metering will flatter revenue in the near term and expose churn in the medium term, because a monthly consumption bill gets reviewed in a way an annual seat count does not.<sup>[13]</sup>

**ASSESSMENT** The structural risk is not that Microsoft loses to a low-code competitor. It is that a meaningful share of the workloads that justified the platform stop being platform-shaped at all.

That does not show up in competitive win rates, which is why it will be visible in net revenue retention before it is visible in any market share analysis.

## For vendors and partners

**ASSESSMENT** Practices built on Power Platform implementation billable hours face the sharper version of this. The defensible position is governance, migration and the Dataverse estate, all of which are durable. The exposed position is building canvas applications for clients at day rates, which is the work most directly substitutable by the client's own staff with an agent.

## What this document cannot answer

Four questions need primary research and are worth commissioning before any large migration decision.

**ASSESSMENT** First, the controlled comparison. Same specification, same operator profile, both paths, wall-clock time to production, repeated enough times to be more than an anecdote. This is the single highest-value study nobody has run, and it would settle the argument in either direction.

**ASSESSMENT** Second, the shadow backlog. What share of internal application demand is refused on Power Platform capability grounds rather than priority grounds? Centers of excellence hold this data and none of it is public.

**ASSESSMENT** Third, real negotiated pricing. Every cost figure in this report uses published rates. Actual Enterprise Agreement terms are the difference between the arithmetic in section six holding and collapsing, and they are not disclosed.

**ASSESSMENT** Fourth, maintenance cost over time. Both paths have a maintainability story and neither has a longitudinal one. Nobody has tracked total cost of ownership of an agent-authored internal application against a canvas application over five years, because the first category is not yet five years old.

## Half-life of this assessment

**ASSESSMENT** The licensing figures in sections four and six decay fastest, within six months. Microsoft has changed Power Platform pricing repeatedly through 2025 and 2026, and any specific rate quoted here should be re-checked against the public pricing page before it is relied on.

**ASSESSMENT** The productivity and code-quality evidence has a twelve-month half-life. Model capability moves quickly, METR has already retired its own study design, and the security and duplication findings are measurements of a moving target.

**ASSESSMENT** The addressable-surface finding is architectural and will hold for the assessment window. Power Platform producing signed desktop binaries or self-hosted container images would require a change in what the product is, not a feature release. If that changes, this report's central argument goes with it.

**ASSESSMENT** The governance distinction between encoding intent and enforcing it is also architectural. It will hold until something in the agent tooling stack can enforce policy at runtime the way a managed environment does, which is a plausible development and would be the most important thing to watch in this whole analysis.

## Limitations

This analysis was not commissioned, reviewed, or funded by Microsoft or by any vendor named in it, and the author holds no position in any of them. No vendor briefings were taken.

The author has worked with Power Platform since 2016 and has built the AI-assisted portfolio described in section five. That is both the source of the field observation and an obvious source of bias, and the reader should weigh the ASSESS-labeled claims accordingly. The FACT-labeled claims are checkable against the cited sources without reference to the author's judgment, and readers who distrust the argument should start there.

The evidence base has three specific weaknesses. The productivity comparison at the center of the question has never been measured by anyone. The code-quality evidence comes predominantly from firms selling code-quality and security tooling, which have a commercial interest in the finding. And the pricing analysis uses published list rates in a market where large buyers do not pay list.

One structural asymmetry is worth stating plainly. Microsoft publishes a great deal about Power Platform and almost nothing about churn, retention, or the share of licensed seats actually in use. Every argument in this report about how the platform is performing commercially is therefore inference from pricing behavior, which is a weaker instrument than the data Microsoft holds and does not release.

## Sources

- [1] Microsoft. "End of AI Builder credits." Microsoft Learn, updated 14 May 2026. *Vendor primary* [learn.microsoft.com](https://learn.microsoft.com)
- [2] Microsoft. "Power Apps pricing." Accessed August 2026. *Vendor primary* [microsoft.com](https://microsoft.com)
- [3] Microsoft. "Power Pages pricing." Accessed August 2026. *Vendor primary* [microsoft.com](https://microsoft.com)
- [4] SAMexpert. "Power Apps Per App: EA and CSP Keep It, MPSA Loses It." 2026. *Practitioner review* [samexpert.com](https://samexpert.com)
- [5] The Licensing Guide. "Power Apps per app plan end of sale." 2026. *Practitioner review* [licensing.guide](https://licensing.guide)
- [6] METR. "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity." 10 July 2025. *Independent research* [metr.org](https://metr.org)
- [7] METR. "We are Changing our Developer Productivity Experiment Design." 24 February 2026. *Independent research* [metr.org](https://metr.org)
- [8] GitClear. "The Maintainability Gap: 2026 AI Code Quality Research." *Third-party comparison (commercial interest possible)* [gitclear.com](https://gitclear.com)
- [9] Veracode. "2025 GenAI Code Security Report." October 2025. *Third-party comparison (commercial interest possible)* [veracode.com](https://veracode.com)
- [10] Veracode. "Spring 2026 GenAI Code Security Update." *Third-party comparison (commercial interest possible)* [veracode.com](https://veracode.com)
- [11] VentureBeat. "Anthropic launches enterprise Agent Skills and opens the standard." December 2025. *Independent press* [venturebeat.com](https://venturebeat.com)

- [12] Agent Skills Directory. Install counts and supported agents. Accessed August 2026. *Practitioner aggregate (self-selected)* [skills.sh](https://skills.sh)
- [13] Vlad Talks Tech. "Power Platform Reborn in the AI Era: Top PPCC25 Keynote Announcements." October 2025. *Practitioner review* [vladtalkstech.com](https://vladtalkstech.com)
- [14] Microsoft. "Pay-as-you-go overview." Microsoft Learn. *Vendor primary* [learn.microsoft.com](https://learn.microsoft.com)
- [15] Soden, David. Published product portfolio. Accessed August 2026. *Practitioner primary* [davidsoden.com](https://davidsoden.com)
- [16] BMad Code Org. "BMAD-METHOD: Breakthrough Method for Agile AI Driven Development." GitHub. *Practitioner aggregate (self-selected)* [github.com](https://github.com)
- [17] InfoWorld. "AI coding tools can slow down seasoned developers by 19%." 2025. *Independent press* [infoworld.com](https://infoworld.com)

## Rights and permitted use

© 2026 David Soden, [davidsoden.com/market-assessment](https://davidsoden.com/market-assessment). All rights reserved.

This document, including its structure, analysis, findings, evidence classification, and framing, is the original work of David Soden, published at [davidsoden.com/market-assessment](https://davidsoden.com/market-assessment).

It is published for public reading. You are free to share the complete, unmodified file, to link to it, and to quote short passages in commentary, reporting, or discussion, provided the author and the site above are credited.

The following require prior written consent: republishing this work in whole or in substantial part under another name, masthead, or brand; excerpting or modifying it in a way that alters the analysis or removes the evidence classifications; incorporating any portion of it into a paid, commissioned, or client-facing deliverable; resale or distribution behind a paywall; and use of this document as training data for machine learning models.

Commissioned Market Assessment reports, commercial licensing, and briefings on this material are available.

This document is analysis, not investment, legal, or procurement advice. It was not commissioned, reviewed, or funded by any company named in it, and the author holds no position in any of them.

Market Assessment reports, commissioned research, and briefings: [davidsoden.com/market-assessment](https://davidsoden.com/market-assessment)