Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Can You Price AI Usage? A Documentation Audit of Credit-Based AI Products

Pratik Gupta

pratikgupta1088@gmail.com

Stripe https://orcid.org/0009-0004-2590-1443

Research Article

Keywords: AI pricing, Credit-based pricing, Pricing transparency, Documentation audit, Usage-based pricing

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11707201/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare potential competing interests as follows: The author is an employee of Stripe. This research was conducted in a personal capacity; the views expressed are the author's own and do not represent Stripe. Stripe was excluded from the study's coded products.

Can You Price AI Usage?

A Documentation Audit of Credit-Based AI Products

Pratik Gupta*

Stripe.

*Corresponding author(s). E-mail(s): pratikgupta1088@gmail.com;

Abstract

Many AI products sell a named credit that users spend as they run actions. Whether a buyer can judge cost depends on what the public documentation shows: how credits map to actions, what a credit costs, and how spending is attributed afterward. This study audits first-party documentation for 22 AI products and asks three questions. Before a run, can a buyer work out the credit cost of a named action? Can the buyer turn those credits into dollars on the same pages? After use, do the pages show where the credits went? The main finding is that documentation is stronger at stating a credit number than at supporting a decision. 17 of 22 products publish an exact credit amount for at least one action. Of these, 7 also state a dollar rate for that action, and 5 of those also have a known or buyer-set quantity, so only 5 reach a computable dollar quote for one event (6 if a rate derived from a same-page plan price counts). On coded pages the usual gap is the dollar rate: 9 products have a known quantity but state none there. A dated read of other same-domain pages (Section 2) raises the products with a dollar rate or a plan price with a credit allowance from 11 to 19 of 22; the added figures are plan prices, not event quotes. Spend after use is less granular still: only 5 of 25 surface records describe an amount per action, job, query or session, and most show totals or a balance only. A credit can carry different prices by route, so a figure needs its route. Providers could close common gaps by stating the dollar value of a credit beside each credit amount, publishing an action-level rate table, and showing credits per action in the usage record. These are implications, not measured effects. The results describe public documentation for a frame of 76 candidates, not product behavior, billing accuracy or market prevalence.

Keywords: AI pricing, Credit-based pricing, Pricing transparency, Documentation audit, Usage-based pricing

---

Independent research conducted in a personal capacity; the views expressed are my own and do not represent Stripe.

1 Introduction

A dollar price per request is directly comparable. A credit price needs two facts: how many credits an action burns, and what a credit costs. This study centers on the first fact and on the visibility of spending afterward. Price per credit is context, used only to test whether the action to credits to dollars chain can be completed from public pages.

The study asks three research questions about what the public pages let a buyer know.

RQ1, before a run. Can a buyer work out from the pages what a named action will cost in credits?

RQ2, credits to dollars. Can the buyer then turn those credits into dollars on the same pages, or is a step missing?

RQ3, after use. Do the pages show where the credits went, and at what level of detail?

The answers come from first-party public documentation only, for 22 AI products, and describe documentation. They do not test product behavior or billing accuracy. Section 2 answers the three questions in order, then gives the field implications, and Section 3 defines the measures behind each question.

2 Main results

Main outcomes. Four outcomes stand out, each measured over the pages read for the 22 primary products.

First (RQ1, before a run), what a buyer can know before a run depends on the action. Of 86 rows, 25 have a quantity fixed per call or event and 18 have a quantity the buyer sets, so 43 can in principle be sized before a run, while 32 depend on run-time quantities such as tokens, compute hours or variable size, and 11 do not apply. Amounts are also uneven inside products: 14 of 22 products have an amount on every recorded row but only 6 have an exact amount on every recorded row, and 11 products mix exact and non-exact rows.

Second (RQ2, credits to dollars), a number is usually published but a complete path to dollars usually is not. 17 of 22 products give an exact credit amount for at least one action. Of these 17, 7 also state a dollar rate on the same pages for that action. For 5 of the 7 the quantity is also fixed or set by the buyer, so a quote for one event is computable (4 with a fixed quantity, 1 once the buyer sets a unit count). The other 2 (Cloudflare Workers AI and GitHub Copilot) price by tokens or neurons, so the dollar rate is known but the cost of an event depends on run-time use. These quotes use a dollar rate stated on the same pages. If a rate derived from a same-page plan price also counts, 6 products reach a quote, with Runway added under a named plan and billing route. 9 products state exact credits and a known quantity but no dollar rate on the same surface, so the last step from credits to dollars cannot be completed from those pages. A dated supplementary read on Nov 6, 2026 covered other pages of the same domains for all 11 products with no stated or plan-derived rate on the coded pages (Adobe Firefly, Clay, Databricks, Gamma, Genspark, Glean, Krea, Listen Labs, Lovable, Suno, Synthesia), and Appendix B3 lists each result.

Counted by rate level (Table R): R1, 9 of 22 products; R2, 11 of 22; R3, 19 of 22. Databricks is not counted: its only dollar rate on those pages is a worked per-DBU example on a different product page (Foundation Model Serving), and the coded GPU serving row states that a rate from another product page is not applied. Counting that example would make R3 20 of 22. Glean and Listen Labs showed no dollar figure in the text read, which does not show that none exists. Listen Labs' own user guide says its platform fee is priced by outlined commitment, and its interview credits have no dollar rate in the text read; Glean's pricing page offers a demo and says nothing about how credits are priced, so Glean is not classed as contract-priced. R3 is rate availability and not a one-event quote: the figures are plan prices with a credit allowance that are specific to a plan and billing route, often bundle other features, sit on different pages from the credit amounts, and the supplementary cohort was not recoded. The quote counts of 5 of 22 and 6 of 22 are unchanged. Across rows the same test is stricter: 10 of 86 rows (12%) have an exact amount, a same-page dollar rate and a known or buyer-set quantity, and by product the counts are Black Forest Labs 2 of 3 rows, HeyGen 3 of 3, Salesforce Agentforce 1 of 1, Snowflake 1 of 7 and World Labs 3 of 8. 7 products reach only an approximate amount, a range or an unknown quantity (4 of them with a dollar rate on the same page, 3 without), and 1 shows no credit amount at all.

Table R. Rate levels, counted separately. Each level is rate availability, not a one-event quote.

Level	What counts	Products	Count of 22
R1	A dollar rate stated for the credit on the coded pages	Black Forest Labs, Cloudflare Workers AI, GitHub Copilot, HeyGen, Notion, Perplexity, Salesforce Agentforce, Snowflake, World Labs	9
R2	R1 plus a rate derived from a plan price on the coded pages	Adds ElevenLabs and Runway	11
R3	R2 plus a plan price with an allowance in the same credit on other pages of the same domain	Adds Adobe Firefly, Clay, Gamma, Genspark, Krea, Lovable, Suno and Synthesia	19
Not counted	No price in the same credit found in the text read	Databricks (a worked example for a different product page), Glean and Listen Labs	3

Third (RQ3, after use), documentation of spend after use is less granular than documentation of price before use. Only 5 of 25 surface records (Gamma, Lovable, Perplexity, Snowflake AI credits and World Labs) describe a per-action, per-job, per-query or per-session amount. 10 describe totals only, 6 a balance only, 1 a log whose credit amount the pages do not describe, and 3 could not be fully captured. A buyer on most surfaces can see how much was spent in a period or in total, but the pages do not show which action spent it. Granularity is the measure here, not usefulness: totals by project, user or team can serve chargeback better than per-action records.

Fourth (context, not a research question), where both a plan price and a top-up price are public, they can imply different credit prices. In the two usable pairs, the plan-implied price per credit is about 1.99 times the top-up price for HeyGen Business and about 0.10 to 0.14 times the top-up price per Action for Relevance AI Pro (an upper bound, because the plan bundles other items). A single credit therefore does not have one price, and a buyer needs to know which route a figure refers to.

These are findings about public documentation. Their practical meaning is a reading of the measured gaps: a buyer who wants to estimate cost before a run finds a fixed or buyer-set quantity documented for about half of the named actions (43 of 86 rows), which is a test of quantity and not of price, while 32 rows depend on run-time quantities, and a buyer who wants to audit spend afterward usually has only totals. Documentation that states the dollar rate on the same page as the credit amount, and shows a per-action amount in the usage record, would close the two most common gaps in this sample. For 8 of the 11 products without a same-page rate, a dated supplementary read found a plan price with a credit allowance on other pages of the same site, so there the gap is placement and route and not total absence of a rate.

What this means for the field. These are implications of the measured gaps, not measured effects. The study does not test whether closing the gaps changes buyer outcomes, support load or revenue. A provider that wants buyers to estimate cost before a run could state the dollar value of a credit on the same page as each credit amount, and say which plan or top-up route the value refers to. It could publish an action-level rate table with the unit and the basis of the quantity, and give a stated range with its assumptions where cost depends on run-time quantities such as tokens. After use, it could provide a usage record that shows credits per action, job, query or run, and document the attribution field as always returned and stable, with a documented, versioned schema. For buyers and researchers the reading is narrower: public credit documentation usually states a credit rate for at least one action, rarely supports a dollar cost for a task, and rarely supports an audit of spend by action. The Discussion sets out each proposal with the measured gap behind it.

3 Measures and coding definitions

How the questions map to the measures. RQ1 uses the row codes, precision and quantity fields. RQ2 uses the matched chain status. RQ3 uses the attribution level. Price per credit is secondary context and is not a research question. The detailed results that follow use these definitions.

Unit of coding and row rule. One row per named action family on one surface. A surface is the plan or interface a page describes. Separate features that the page prices separately get separate rows (for example Glean search, assistant, meeting notes). Model or size variants of one action type are collapsed into one row when they share one precision class (for example fixed-price video models). A variant with a variable or from-price keeps its own row. Worked examples are not counted when the underlying rates are coded, and rate multipliers are not rows. Rows that appear on out-of-scope or out-of-credit products are not counted. Row counts depend on this rule and on how pages are organized, so they are secondary to the product tiers. A sensitivity analysis drops the two products with more than 8 rows.

RQ1 (ex ante). Code: D, a fixed number per named unit; C, a range, worked example, average or formula; P, unit named but no amount, or a relative multiplier with no baseline; N, the page says cost cannot be known in advance; I, capture incomplete. Precision: exact (no qualifier), approximate (the page says approximate, such as ElevenLabs and the ChatGPT rate card), range or estimate. Quantity known before

run: yes (fixed per call or event), user_set (a unit count the buyer chooses, such as seconds), or no (depends on run-time quantities such as tokens, compute hours, neurons or variable size).

Existence flag (not a transparency ranking). T1: at least one exact numeric unit rate observed for a named action. T2: numeric amounts observed but none exact. T3: no numeric amount observed on the pages read. The flag shows only that a number exists. It says nothing about how many of a product's actions are priced, whether quantities are known before a run, or whether a zero or promotional figure informs burn. Coverage is reported separately (Section 4).

RQ2 (credits to dollars): matched chain status (per row). C1 exact credits, fixed quantity per event, and a same-surface dollar rate on the same pages. C1u as C1 but the quantity is a unit count the buyer sets. C2 exact credits and known quantity, no same-surface dollar rate. C3 approximate, range or unknown quantity, with a same-surface rate. C4 approximate, range or unknown quantity, no rate. C0 no credit amount. A rate from another surface is never used. A dollar rate counts when the pages of the one coded surface that carry the credit amount also state a price per credit or per credit pack for a named route, such as pay-as-you-go, top-up or list price. A rate derived by dividing a plan price by its credit allowance is not counted in the main results. Section 4 reports it as a sensitivity. Chain status describes how complete a quote for one event is. It combines a stated dollar rate with a quantity known before the run, so it is not a ranking of transparency, and Section 4 also reports the two conditions separately.

RQ3 (ex post). One record per product and surface with an attribution level: A per action, request, job, query or usage event; S per session, run or chat; T totals only (period, feature, model, project, user, hour window); B balance only; U a usage log is described but the pages read do not say it shows a credit amount; I not captured. A level below A or S means the pages read do not describe a finer record. It does not mean no per-action log exists. RQ3 measures documented attribution granularity. It does not test whether a usage record is complete, consistently returned or reconcilable with an invoice, and it is not a test of billing correctness.

Secondary context (H1, not tested). The focus is transparency, and price per credit is context. A subscription may carry a lower plan-implied price per credit than a top-up. For this contrast the study records bounded pairs only where the same product pages state a plan price, the units that plan includes and a top-up price. The plan-implied price is the plan price divided by included units. Plans bundle seats and features, so it is not the price of the unit alone. Two products gave a usable pair. Two products cannot support a hypothesis test, so H1 is illustrated and not tested.

4 Detailed results

Existence flag (primary population, n = 22). T1 17, T2 4, T3 1. Table 1 lists each primary product with its dataset row, surface and results. Table 3 shows other frames and exclusions, and Table 4 lists the broadened surfaces apart.

RQ1: rows and coverage (primary population). 86 counted rows: D 46, C 28, P 8, N 1, I 3. Precision: exact 40, approximate 6, range or estimate 28, none 12.

Quantity: 25 yes, 18 user-set, 32 no, 11 not applicable. Coverage is measured only over the action families recorded on the pages read. It is not a census of each product's actions, and families the pages do not name are invisible to it. Over recorded rows, 14 of 22 products have an amount (D or C) on every recorded row, and 6 have an exact amount on every recorded row. 11 products mix exact and non-exact rows. The one zero or promotional row is tested in Table 3.

RQ2: matched chains (primary population). Rows: C1 6, C1u 4, C2 20, C3 18, C4 26, C0 12. Best status per product: C1 4, C1u 1, C2 9, C3 4, C4 3, C0 1. Same-action classification of the 17 products with an exact credit amount, with the dollar rate and the quantity judged on the same action row: 5 have both (Black Forest Labs, HeyGen, Salesforce Agentforce, Snowflake AI credits, World Labs); 2 have a dollar rate but a run-time quantity on every such action (Cloudflare Workers AI, GitHub Copilot); 9 have a known quantity and no dollar rate on any exact action (Adobe Firefly, Clay, Glean, Krea, Listen Labs, Lovable, Runway, Suno, Synthesia); 1 has neither (Databricks). The other 5 products have no exact credit amount on any action (ElevenLabs, Gamma, Genspark, Notion, Perplexity). If a plan-derived rate counts, Runway moves to the first group. Across rows, C1 and C1u are 10 of 86 (12%), and by product the counts are Black Forest Labs 2 of 3 rows, HeyGen 3 of 3, Salesforce Agentforce 1 of 1, Snowflake 1 of 7 and World Labs 3 of 8. The best-status figures describe how complete a one-event quote is, not a ranking of transparency. The C1 rows are: HeyGen (AI model training); Salesforce Agentforce (Agentforce Action); Black Forest Labs (FLUX.1 family image generation and editing (fixed credits per image: Kontext, 1.1, Fill)); World Labs (World generation usage event (marble-1.0, marble-1.1 fixed credits; draft)); World Labs (Frameboy session in Marble Studio (up to one minute)); World Labs (HQ mesh export (Export API)). Each is a quote for one event, unit or component. Table 2 shows the worked chains, with the amount and peg quoted and the condition on what the dollar figure covers. For example, the World API figure covers one world generation event and a panorama event may be billed separately, and the Relevance AI figure is the Action component only, with separately billed Vendor Credits (a product term) excluded. World Labs World API has a $1.00 per 1,250 credits rate. The Marble app pages state no dollar rate and the Marble 2 Developer API beta page states a different rate, a separate surface that is not equated. Cloudflare Workers AI mixes units across its image and audio families, so it is not a C1 row. Pegs are public list prices for a named route, and negotiated discounts were not observed. Snowflake's rate depends on the route: $2.00 per AI Credit on demand globally, $2.20 regional, and $1.88 to $2.00 on capacity tiers (a discount of up to about 6% at larger annual commitments). Sensitivity to the rate rule: counting a rate derived from a same-page plan price divided by its credit allowance would add Runway, whose derived price per credit differs by plan and billing route, and would move ElevenLabs from C4 to C3. Under that rule the rows are C1 6, C1u 5, C2 19, C3 25, C4 19, C0 12, the best status per product is C1 4, C1u 2, C2 8, C3 5, C4 2, C0 1, and 6 of 22 products have a computable quote under a named route.

Sensitivity. Table 3 shows the flag for the primary population, with the different-product replacements added (Datadog Bits AI, Devin, Intercom Operator), with the two products that have more than 8 rows dropped, and with zero or promotional rows

dropped, and then the same exclusions on the broadened frame. These are alternative frames and exclusions, not an alternative extraction rule. 3 primary products have an incomplete-capture row (Adobe Firefly, Databricks, Krea) and 4 in the broadened frame (Adobe Firefly, Databricks, Krea, Microsoft 365 Copilot). Snowflake was first coded I because its AI-function rate table (Table 6) was cut off in the capture; it was recoded from the full official PDF (effective Oct 29, 2026). The Snowflake rows are a selected, non-exhaustive subset of Table 6 (one model per function type plus named non-token features), because the table lists many models and Platform Credit features, so recorded-row coverage for Snowflake is not completeness. The results are conditional on the captured records. Unvisited pages, sources not located and membership gaps are not bounded here, so no real upper bound is given and no claim is made about whether truncation explains any T3 result.

RQ3 (primary population, n = 25 surface records for 22 products). A 4, S 1, T 10, B 6, U 1, I 3. Products with at least one A or S surface: Gamma, Lovable, Perplexity, Snowflake (AI credits), World Labs. Caveats from the pages: Gamma reports credits per job only in the API and says the field is informational and may be omitted; Snowflake attributes per query, function, model and warehouse inside one-hour windows. The Genspark log is U because the pages do not say each entry shows a credit amount. A balance or totals view in the pages read does not imply that no per-action logs exist. In the broadened frame (32 records): A 4, S 4, T 14, B 6, U 1, I 3. The added S records are ChatGPT lifetime credits per Work or Codex chat (stated as not a complete record) and the Microsoft Cowork /cost command, which shows approximate credits for the open task and states it is not a billing record; Microsoft also documents an admin Cost Management view by user, group, agent and service (T).

Table 1. Primary population, one line per product with the actual product first: dataset row, surface coded, existence flag, RQ1 rows with counts, best chain (RQ2), RQ3 levels.

- **Adobe Firefly**. Dataset row: Adobe Firefly. Surface coded: Adobe Firefly generative credits (Firefly FAQ page). Flag T1; rows 11 (exact 9; D 9, C 0, P/N/I 2); RQ3 B; best chain C2.
- **Black Forest Labs**. Dataset row: Black Forest Labs. Surface coded: Black Forest Labs API and Playground credits. Flag T1; rows 3 (exact 2; D 2, C 1, P/N/I 0); RQ3 T; best chain C1.
- **Clay**. Dataset row: Clay. Surface coded: Clay Data Credits. Flag T1; rows 3 (exact 2; D 2, C 1, P/N/I 0); RQ3 T; best chain C2.
- **Cloudflare Workers AI**. Dataset row: Cloudflare Workers AI. Surface coded: Cloudflare Workers AI (Neurons). Flag T1; rows 2 (exact 2; D 2, C 0, P/N/I 0); RQ3 T; best chain C3.
- **Databricks**. Dataset row: Databricks. Surface coded: Databricks Model Serving pricing page (GPU serving); Databricks pricing page. Flag T1; rows 2 (exact 1; D 1, C 0, P/N/I 1); RQ3 I; best chain C4.
- **ElevenLabs**. Dataset row: ElevenLabs. Surface coded: ElevenLabs web app credits. Flag T2; rows 7 (exact 0; D 6, C 1, P/N/I 0); RQ3 T; best chain C4.

- **Gamma**. Dataset row: Gamma. Surface coded: Gamma API (developers.gamma.app); Gamma app (help center). Flag T2; rows 4 (exact 0; D 0, C 2, P/N/I 2); RQ3 B,A; best chain C4.
- **Genspark**. Dataset row: Genspark. Surface coded: Genspark credits guide (help center). Flag T3; rows 1 (exact 0; D 0, C 0, P/N/I 1); RQ3 U; best chain C0.
- **GitHub Copilot**. Dataset row: GitHub Copilot. Surface coded: GitHub Copilot AI Credits (usage-based, over included allowance). Flag T1; rows 2 (exact 1; D 1, C 0, P/N/I 1); RQ3 T; best chain C3.
- **Glean**. Dataset row: Glean. Surface coded: Glean Enterprise Flex (FlexCredits). Flag T1; rows 11 (exact 1; D 1, C 10, P/N/I 0); RQ3 T; best chain C2.
- **HeyGen**. Dataset row: HeyGen. Surface coded: HeyGen credit-based plans. Flag T1; rows 3 (exact 3; D 3, C 0, P/N/I 0); RQ3 B; best chain C1.
- **Krea**. Dataset row: Krea. Surface coded: Krea image generation (docs, compute units); Krea other products (video, enhance, edit). Flag T1; rows 3 (exact 1; D 1, C 0, P/N/I 2); RQ3 T; best chain C2.
- **Listen Labs**. Dataset row: Listen Labs. Surface coded: Listen Labs Interview Credits (user guide). Flag T1; rows 4 (exact 1; D 1, C 3, P/N/I 0); RQ3 B; best chain C2.
- **Lovable**. Dataset row: Lovable. Surface coded: Lovable credits. Flag T1; rows 2 (exact 1; D 1, C 1, P/N/I 0); RQ3 A; best chain C2.
- **Notion**. Dataset row: Notion. Surface coded: Notion credits (Agent, Standard and Select allowances). Flag T2; rows 3 (exact 0; D 0, C 3, P/N/I 0); RQ3 T; best chain C3.
- **Perplexity**. Dataset row: Perplexity. Surface coded: Perplexity Computer credits (consumer plans). Flag T2; rows 4 (exact 0; D 0, C 4, P/N/I 0); RQ3 S; best chain C3.
- **Runway**. Dataset row: Runway. Surface coded: Runway web app credits. Flag T1; rows 1 (exact 1; D 1, C 0, P/N/I 0); RQ3 B; best chain C2.
- **Salesforce Agentforce**. Dataset row: Salesforce Agentforce. Surface coded: Salesforce Agentforce Flex Credits (pricing page). Flag T1; rows 1 (exact 1; D 1, C 0, P/N/I 0); RQ3 I; best chain C1.
- **Snowflake (AI credits)**. Dataset row: Snowflake (AI credits). Surface coded: Snowflake Service Consumption Table (AI Credits and Platform Credits). Flag T1; rows 7 (exact 7; D 7, C 0, P/N/I 0); RQ3 A; best chain C1u.
- **Suno**. Dataset row: Suno. Surface coded: Suno credits (help center). Flag T1; rows 3 (exact 1; D 1, C 0, P/N/I 2); RQ3 B; best chain C2.
- **Synthesia**. Dataset row: Synthesia. Surface coded: Synthesia self-serve credits. Flag T1; rows 1 (exact 1; D 1, C 0, P/N/I 0); RQ3 T; best chain C2.
- **World Labs**. Dataset row: World Labs. Surface coded: World Labs Marble (docs.worldlabs.ai/marble/models); World API (docs.worldlabs.ai API billing); Marble 2 Developer API (beta billing page); Marble Studio (beta billing page). Flag T1; rows 8 (exact 5; D 5, C 2, P/N/I 1); RQ3 A,I,T; best chain C1.

Table 2. Worked chains (units x same-surface dollar rate), each for one stated event or unit and not for a whole task.

- **Black Forest Labs**, FLUX.1 Kontext [pro] image: 4 credit x $0.01 per credit = $0.0400. Amount quote (line 47): "| FLUX.1 Kontext [pro] | 4 credits | $0.04/image |" https://help.bfl.ai/articles/7986977817-what-are-the-costs-associated-with-using-your-models. Peg quote (line 12): "1 credit = $0.01 USD" https://help.bfl.ai/articles/7986977817-what-are-the-costs-associated-with-using-your-models. Condition: Fixed per-image credit for this model only; FLUX.2 is priced per megapixel and is not covered.
- **Black Forest Labs**, FLUX1.1 [pro] Ultra image: 6 credit x $0.01 per credit = $0.0600. Amount quote (line 55): "| FLUX1.1 [pro] Ultra | 6 credits | $0.06/image |" https://help.bfl.ai/articles/7986977817-what-are-the-costs-associated-with-using-your-models. Peg quote (line 12): "1 credit = $0.01 USD" https://help.bfl.ai/articles/7986977817-what-are-the-costs-associated-with-using-your-models. Condition: Fixed per-image credit for this model only; FLUX.2 is priced per megapixel and is not covered.
- **Salesforce Agentforce**, Agentforce Action: 20 credit x $0.005 per credit = $0.1000. Amount quote (line 156): "Agentforce actions are 20 Flex Credits" https://www.salesforce.com/agentforce/pricing/. Peg quote (line 34): "$ 500 USD/Per 100k Credits" https://www.salesforce.com/agentforce/pricing/. Condition: Per action; a use case has several actions and the number run is not known before the run. Public list price (US, $500 per 100,000 credits); negotiated discounts were not observed.
- **Snowflake (AI credits)**, AI_PARSE_DOCUMENT Layout, 1,000 pages: 3.66 credit x $2.00 per credit = $7.3200. Amount quote (line 1726): "3.66 AI Credits per 1,000 pages" https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf. Peg quote (full capture, line 417): "Global $2.00" https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf. Condition: Per 1,000 pages; the page count is set by the buyer. On-demand global rate; regional $2.20; capacity tiers $1.88 to $2.00, with discounts starting at $1.2 million of annual commitment (tier ranges shown as illustrative). Public list price; negotiated discounts were not observed.
- **Microsoft Copilot Credits (Work IQ Tools API, Copilot Credits API name, not Microsoft 365 Copilot seats)** [secondary, not a primary-population chain], Work IQ Tools API call: 0.1 credit x $0.01 per credit = $0.0010. Amount quote (line 347): "0.1 Copilot Credit per API call for actions and tools." https://www.microsoft.com/licensing/guidance/CopilotCreditsGuide. Peg quote (line 90): "Pricing: $0.01/Copilot Credit" https://www.microsoft.com/licensing/guidance/CopilotCreditsGuide. Condition: One API call only, not the whole task; calls per task are not stated. Pay-as-you-go rate; the Pre-Purchase Plan (P3) tiers give a 5% to 20% discount, and unused credits expire at the end of the annual term.
- **HeyGen**, Avatar IV photo look, 1 minute: 16 credit x $0.05 per credit = $0.8000. Amount quote (line 54): "Avatar IV Photo Look: 16 credits per minute" https://help.heygen.com/en/articles/15126059-how-to-use-credits-on-heygen. Peg quote (line 41): "Additional Business credits cost $0.05 per credit" https://help.heygen.com/en/articles/15126059-how-to-use-credits-on-heygen. Condition: One minute of Avatar IV video; video length is set by the user; Business extra-credit rate. Plan-derived comparison (a sensitivity, not a stated rate): the Business plan at $149 per month for 1,500 credits is $0.099 per credit, so the same

minute would be about $1.59; the plan bundles seats and features, so this is an upper bound on the price of the credits.

- **World Labs (World API)**, marble-1.1 world generation: 1500 credit x $0.0008 per credit = $1.2000. Amount quote (line 51): "the world generation usage event is billed at 1,500 credits" https://docs.worldlabs.ai/api/pricing. Peg quote (line 17): "$1.00 USD per 1,250 credits" https://docs.worldlabs.ai/api/pricing. Condition: World generation event only; a panorama event may be billed separately.
- **Perplexity**, Quick fact or definition (low end, 15 credits): 15 credit x $0.01 per credit = $0.1500. Amount quote (line 13): "A quick fact or definition uses about 15 to 70 credits" https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity. Peg quote (line 13): "100 credits equals $1" https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity. Condition: Page gives a range with examples, not a guarantee; endpoints of that range shown.
- **Perplexity**, Quick fact or definition (high end, 70 credits): 70 credit x $0.01 per credit = $0.7000. Amount quote (line 13): "A quick fact or definition uses about 15 to 70 credits" https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity. Peg quote (line 13): "100 credits equals $1" https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity. Condition: Page gives a range with examples, not a guarantee; endpoints of that range shown.
- **Relevance AI** [secondary, not a primary-population chain], One Action (Tool run): 1 Action x $0.08 per Action = $0.0800. Amount quote (line 13): "An Action is a single run of a Tool." https://relevanceai.com/docs/admin/subscriptions/plans. Peg quote (line 33): "Actions: $80 USD per 1,000 Actions" https://relevanceai.com/docs/admin/subscriptions/plans. Condition: Action credits only; separately billed Vendor Credits (product term) for model use are excluded, so this is not the total tool cost.

Table 3. Existence flag under alternative frames and exclusions (T1 / T2 / T3)				
Frame	Products	T1	T2	T3
Primary population: dataset row with exact coded surface	22	17	4	1
Primary, dropping Adobe Firefly and Glean (more than 8 rows)	20	15	4	1
Primary, dropping 1 zero or promotional row	22	17	4	1
Primary plus different-product replacements (secondary)	25	18	4	3
Broadened family frame, all 28 coded products (secondary)	28	20	5	3
Broadened, dropping Adobe Firefly and Glean	26	18	5	3
Broadened, dropping 2 zero or promotional rows	28	19	6	3

Table 4. Broadened family frame, surfaces that are not the original dataset row, under their real names. Counts here are not part of the primary results. Broadened rows: 108 (D 61, C 29, P 13, N 1, I 4); precision exact 44, approximate 17, range or estimate 29, none 18.

- **ChatGPT Business, Enterprise and Edu credit rate card (OpenAI)**. Dataset row: OpenAI (ChatGPT consumer). Surface coded: ChatGPT credit-based workspaces: Business/Enterprise/Edu rate card (help article); NOT the consumer

plans. Flag T2; rows 6 (exact 0; D 6, C 0, P/N/I 0); RQ3 S; best chain C4. Membership: surface differs from dataset row.

- **Datadog Bits AI**. Dataset row: Datadog (LLM Observability). Surface coded: Datadog AI credits (pricing page). Flag T1; rows 6 (exact 1; D 6, C 0, P/N/I 0); RQ3 T; best chain C2. Membership: different product replacing the dataset row.
- **Devin (Cognition)**. Dataset row: Cognition (Devin + Windsurf). Surface coded: Devin billing (docs). Flag T3; rows 2 (exact 0; D 0, C 0, P/N/I 2); RQ3 T; best chain C0. Membership: different product replacing the dataset row.
- **Intercom Operator**. Dataset row: Intercom Fin. Surface coded: Intercom Operator credits (Pro add-on). Flag T3; rows 1 (exact 0; D 0, C 0, P/N/I 1); RQ3 T; best chain C0. Membership: different product replacing the dataset row.
- **Microsoft Copilot Credits for Copilot Studio and usage-based features**. Dataset row: Microsoft 365 Copilot. Surface coded: Microsoft Copilot Credits (licensing guide; dataset row "Microsoft 365 Copilot"). Flag T1; rows 5 (exact 2; D 2, C 1, P/N/I 2); RQ3 S,T; best chain C1. Membership: surface differs from dataset row.
- **Relevance AI**. Dataset row: not in the dataset. Surface coded: Relevance AI Actions; Relevance AI Vendor Credits (product term). Flag T1; rows 2 (exact 1; D 1, C 0, P/N/I 1); RQ3 S; best chain C1. Membership: added, not in dataset.

H1 context pairs (secondary, not tested). HeyGen Business: $149 per month for 1500 credits, plan-implied $0.099 per credit, against a top-up of $0.05 per credit, which is about 1.99 times the top-up. The plan bundles seats and features. HeyGen Creator: $29 for 600 credits, $0.048 per credit, and the page says Creator has no one-time purchases, so there is no pair. Relevance AI Pro: $29 per month on monthly billing, or $19 on annual billing, for 2500 Actions. Plan-implied price per Action is $0.0116 on monthly billing and $0.0076 on annual billing, against a top-up of $0.08 per Action, which is about 0.14 and 0.10 times the top-up. These are upper bounds because the plan also includes Vendor Credits and build users. The two plan-implied quotients fall on opposite sides of the top-up price. Because the bundle allocation is unknown, a quotient above the top-up price (HeyGen Business) does not show that subscription credits are expensive, and a quotient below it (Relevance AI) is an upper bound. These are illustrative arithmetic, not standalone discounts. No other product was paired. Subscription discounts are not established here.

The per-product inventory of recorded rows by precision and quantity is not reproduced here.

5 Discussion

This study asks two things: whether public documentation lets a buyer anticipate the credit cost of an action before it runs, and whether it lets them see the spend afterward. What the documentation read supports: for a named action or unit, exact credit amounts are public at 17 of 22 primary products, but the flag says nothing about how many actions are covered, and 32 of 86 rows depend on run-time quantities such as tokens, compute hours or variable size that the buyer cannot fix before a run. A dollar figure for one stated event can be computed from the pages for 6 rows (HeyGen

(AI model training); Salesforce Agentforce (Agentforce Action); Black Forest Labs (FLUX.1 family image generation and editing (fixed credits per image: Kontext, 1.1, Fill)); World Labs (World generation usage event (marble-1.0, marble-1.1 fixed credits; draft)); World Labs (Frameboy session in Marble Studio (up to one minute)); World Labs (HQ mesh export (Export API))), and for 4 more once the buyer sets a unit count. These are quotes for a single event, call, image, action or component, and the conditions in Table 2 state what each covers. None of them predicts the cost of a whole task, because the pages do not say how many events a task triggers. For agent products the pages give ranges, averages and examples at best. What the documentation cannot establish: what logged-in screens or in-product previews show, whether providers could estimate a task cost in advance, and whether any figure matches billing. Where a page says an in-product preview exists, such as Gamma showing image credit cost before generation, it is recorded as a pointer and the amount is not visible. After the run, pages that describe an action-level or job-level record of spending (A or S) are rarer than pages that describe balances or totals. Whether providers should show a pre-run estimate is a separate, tentative question and is not tested here.

5.1 Implications for providers of credit-based products

Observed capabilities. In the pages read, three capabilities appear at some products and are absent at most: an exact credit rate for a named action (17 of 22 products have at least one), a dollar rate for the credit on the same surface as the credit amount (7 of the 17 products with an exact amount state one, and 5 of 22 reach a computable one-event dollar quote), and a usage record that shows the credit amount per action, job, query or session (5 of 25 surface records). Examples in the pages: Gamma returns a per-job credit amount in its API, but the page calls the field informational and says it may be omitted; Snowflake attributes spend per query, function, model and warehouse; and HeyGen, Salesforce Agentforce, Black Forest Labs and World Labs publish a fixed credit amount next to a dollar rate (Table 2). A dated supplementary read of other same-domain pages found a plan price with a credit allowance for 8 more products, so a dollar price somewhere on the domain is more common (19 of 22, as plan prices with credit allowances and not event quotes) than the same-surface count suggests. Products may offer previews and logs inside logged-in screens that these pages do not describe.

Proposed capabilities. The items below are proposals that follow from the gaps measured here. They are not tested, and this study does not measure whether they change buyer behavior, support load or revenue.

1. State the dollar value of a credit on the same page as each credit amount, and say which plan or top-up route the value refers to. Rationale: 9 primary products (8 if a plan-derived rate counts) state exact credits and a known quantity but no same-surface dollar rate, and the two plan and top-up pairs imply different prices per credit. For 6 of the 9, plan prices appear on other pages of the same site in the dated supplementary read, so the proposal concerns placement and the named route.

2. Publish an action-level rate table that covers every priced action, with the unit and the basis of the quantity. Rationale: 11 products mix exact and non-exact rows, and only 6 products have an exact amount on every recorded row.

3. For actions whose quantity depends on run-time values such as tokens, compute hours or variable size, give a pre-run estimate or a stated range with the assumptions behind it. Rationale: 32 of 86 rows depend on run-time quantities, and for agent products the pages give ranges, averages and examples at best.

4. Provide a usage record that shows the credits spent per action, job, query or run, with the action name, time and amount, and not only a balance or a period total. Rationale: 10 of 25 surface records describe totals only and 6 a balance only.

5. Document the attribution field as always returned and stable, and document its schema with versioning, so that a buyer can build reporting on it. Rationale: the one per-job field found in an API is described as informational and possibly omitted.

Why these matter: a buyer needs to estimate cost before a run and to check the spend afterward, and each gap above breaks one step of that path.

Provider-side considerations. The pages give some context for the patterns measured, without showing why any provider chose its design. Several describe amounts that vary with the task: Lovable says credits vary with task complexity in Default Mode, and Glean and Notion publish typical or estimated amounts. That limits what an exact pre-run amount can promise, so proposal 3 asks for a stated range with its assumptions and not a fixed number. Of the 22 products, 17 state a rule on expiry or rollover somewhere in the captured pages (Appendix B4), 19 describe what happens when the allowance runs out (2 of them only conditionally), and 6 say anything about failed runs (3 of them conditionally) (Appendix B5); the rules differ, which affects what a dollar rate per credit means across plans. Whether credits expire, roll over or are charged on retries changes what a plan-implied price per credit is worth to a buyer, but the study records only that such statements exist, not their effect on cost. Publishing a dollar rate, a rate table or a per-action usage record could carry costs for providers, which this study did not examine, so the proposals are not a claim that the benefit to buyers outweighs them.

Route comparison (within product only). Where a page states more than one route for the same credit, the routes differ. The annual-billing price per month is about 44% below monthly for Krea ($5 against $9 for 5,000 units), 20% below for Genspark ($19.99 against $24.99 for 10,000 credits) and for Suno ($8 against $10 for 2,500 credits), about 17% below for Lovable at 100 credits ($250 a year against $25 a month), and about 34% below for Relevance AI Pro ($19 against $29 for 2,500 Actions). Plan-implied prices also differ from top-up prices (HeyGen Business and Relevance AI Pro, Section 2). These comparisons hold the product, the credit and the allowance fixed. No dollars-per-credit comparison across products is given, because a credit is not a common unit and the outputs, features and conditions differ.

6 Related work and positioning

Closest prior material, read from public pages on Nov 4, 2026:

AI Tools Police, Pricing Transparency Audit (Kaya; project page and freeze Aug 17, 2026; version 1.4 released Sep 4, 2026, version DOI 10.5281/zenodo.22303623, all-versions DOI 10.5281/zenodo.22234975). Documents-only audit of 76 consumer AI products against a pre-published instrument. Of 48 products with a credit system, 38

define the credit unit and 20 (41.7%) publish the credit-to-output rate in full, 18 in part and 10 not at all; cost per output is computable for 49 of 72 products (68.1%); 7 of 64 document charging for failed generations. These figures were read directly from the live audit page on Nov 4, 2026. This is the nearest work on RQ1-type disclosure. It is consumer focused and, in the parts read, does not code ex post per-action records [1].

Strategy of Finance, AI Credit Index (edition 1, observed Jun 2026). Reports that 60% of 139 top-up tiers publish the top-up price and that 12 of 52 products peg a credit to a fixed dollar value. Focus is top-up price disclosure and credit-dollar pegs. Its definition of a peg is one credit tied to a fixed dollar value, not a variable amount of compute. Edition 1 page [2]. Edition 2 (Q3 2026), read on the live page on Nov 6, 2026: 75% of 208 top-up tiers publish the price, and 18 of 75 products peg a credit to a fixed dollar value [3]. Edition 2 is a dated supplementary reference and was not part of the Nov 4 comparison.

UsagePricing, credit equation published in halves. Practitioner count of companies that publish both halves of the equation [4].

TrueUpp, 64 usage-based pricing pages (Oct 9, 2026). Codes true-up, rollover and plan terms in usage-based pricing pages, not action-level cost [5].

IVRIS Tech, AI Pricing Transparency source audit (purposive panel of 100 named B2B software companies; version 1.1, Jul 26, 2026, updated Sep 21). Reports the minimum AI-access cost as calculable from public sources for 47, not calculable from retrieved sources for 37 and unresolved for 16; marginal usage cost calculable for 14. Focus is cost calculability, not credit units or ex post records [6].

DataDab, 106-company coding study (Oct 2026). Classifies how B2B SaaS packages AI across 286 pricing pages, including metered credits (32 of 91 priced offers) and the premiums that are computable (three) [7].

Gonzalez, Moving Targets (SSRN 7162638). Documentary analysis of post-purchase changes to pricing and access terms at four providers [8].

PricingIO practitioner note on forecasting credits (Jul 2026) [9].

Token billing audits: Sun et al., Invisible Tokens, Visible Bills (PMLR v306), and Yue et al., Observable Correlates of Token Pricing (ACM DEAI 2026, doi 10.1145/3837873.3837924). These concern token billing, not credit units [10], [11].

An earlier study by the same author on AI pricing models [18]. It characterizes AI pricing models and lists consumption-rate transparency as a risk. The present study follows up on that one dimension.

Spend management and usage-based billing (read Nov 6, 2026). The FinOps Foundation framework describes budgeting, forecasting and allocation as capabilities of the buyer's organization [12], [13], [14]. Its FinOps for AI overview [15] applies these capabilities to AI spend, and its paper on a generative AI cost and usage tracker [16] describes how a buyer can track and attribute token-based cost. In the parts read, these are guidance for buyers on managing spend, not audits of what providers publish about credits. Bala and Carr, Usage-based pricing of software services under competition [17], compare usage-based and fixed pricing for software services in a competitive model. It is a pricing-theory study and does not examine documentation. This audit is narrower: it asks what provider documentation lets a buyer know about credit cost before and after use. It does not study the buyer-side

controls that the FinOps sources describe. Clear provider pricing documentation can help a buyer estimate and review spend, but this study does not test whether it makes spend predictable or controllable.

Coverage and limits of this review. Searches were run on the open web plus site-restricted queries for ACM, SSRN and IEEE (Nov 4, 2026). Result pages were screened by title and, for the items above, by reading the fetched page. The IEEE queries returned no relevant hits in the text returned. Google Scholar, Scopus and Web of Science were not searched. This is a bounded review and not a systematic one, so no novelty claim is made. The angle this study pursues is the ex post dimension (RQ3) with a named attribution level, action-family coding with stated input observability, matched credit-to-dollar chains on the same surface, and per-row evidence with a hash of each captured page (Appendices B1 to B5). Whether any of this is new has not been verified. The 76-row frame in this study is a separate dataset from the 76-product AI Tools Police audit.

Positioning against the closest work. The AI Tools Police audit codes 76 consumer products against a 37-variable instrument with 90 dated deviations on the page as read on Nov 4, 2026. The closest variables to RQ1 are its credit-to-output rate and cost-per-output computability, quoted above from its live page. This study differs in surface (a mix of consumer, enterprise and agent products), in the ex post RQ3 coding, in coding at the action-family grain, and in its same-surface chain test. The headline figures are not directly comparable. The AI Tools Police figure of 49 of 72 products (68.1%) with cost per output computable is measured on its own instrument at the product level. Strategy of Finance counts products that peg a credit to a fixed dollar value (12 of 52 in Edition 1, 18 of 75 in Edition 2). No directly comparable peg count exists here, because peg type was not coded and the frames differ. The nearest stated-dollar-rate count here is 9 of 22 products (41%, or 11 of 22, 50%, if plan-derived rates count). It is a different measure, and it includes rates that vary by route. The 5 of 22 figure for a computable one-event quote is a stricter test that also needs a known quantity (6 of 22 with a plan-derived rate). UsagePricing's credit-equation trend page is marked superseded as of Sep 25, 2026: a by-name recount found 10 companies publishing both halves of the equation and 7 publishing one half. Its list of both-halves companies includes Snowflake Cortex and Databricks Mosaic AI. Here Databricks is coded I for part of its pages, which is a collection gap in this study and not evidence of opacity. Snowflake Table 6 was unreadable with the capture tool at first and is coded from the full official Service Consumption Table PDF.

7 Method

Frame and inclusion. 76 candidates: a 75-row company dataset plus one added product (Relevance AI). The 75-row dataset is the company list of an earlier pricing study by the same author [18]. Which saved version of that list the frame was copied from was not recorded. The 76 candidates and each one's status and basis are listed in Appendix A. The URL, date and hash of every captured page were recorded. A search log for the page-finding step, including candidates screened out, was not kept, so that step cannot be repeated exactly. First-party public pages only. No logins,

accounts, paid tests or contact. Any company affiliated with the author is excluded. Included when the product markets a named unit separate from dollars with a stated or derivable link to actions. A wallet balance in dollars is listed and not coded. Outcome-only pricing is out of scope.

Screening trace. An early row-level screen proved unreliable: 6 of 27 rows first marked out flipped once pages were read. All final statuses are page-based. Counts: 28 in scope (coded); 9 wallet (page-verified), 1 wallet (search excerpt only), 5 wallet (listed, not coded); 16 out of scope (page-verified); 15 with no credit pricing located (a search limit, not a finding); 2 capture incomplete. Appendix A lists all 76. Six coded products are not the original dataset row on the same surface. Consumer ChatGPT and Microsoft 365 Copilot were coded from a different surface (the Business, Enterprise and Edu credit rate card; Copilot Credits for Copilot Studio and usage-based features). Datadog Bits AI, Devin and Intercom Operator are different products or SKUs from their dataset rows (Datadog LLM Observability, Windsurf plus Devin, and Fin, which is priced per resolution), and are not treated as partial versions of those rows. Relevance AI is not in the dataset. These six are reported apart, under their real names.

Coding protocol. Definitions were drafted in pilot coding on Nov 3 and fixed on Nov 4, 2026 at 12:02 AM PT. This is a pilot-informed freeze, not a pre-registration. All rows in this version were coded from scratch at the action grain after the freeze, applying the row rule, the precision field and the quantity field, and after the pilot findings were reviewed. After the freeze, a review on Nov 6, 2026 found that one worked-example row (Synthesia, script edit re-render) had been counted although the row rule excludes worked examples of a coded rate. It was removed, so counted rows changed from 87 to 86 and rows with a known quantity from 44 to 43; the stated-rate count stayed at 5 of 22. The same review added the plan-derived-rate sensitivity, the same-action classification and a Snowflake worked chain, and corrected the Salesforce amount citation to the dated capture. A later check the same day widened the cited line span of one row (Salesforce Agentforce Action) to include the line that states the amount, and moved the Snowflake peg citation to the full capture; no code changed. No other row was recoded.

Supplementary same-domain read (Nov 6, 2026, separate from the frozen Nov 4 baseline). The 11 products with no stated or plan-derived dollar rate on the coded pages (Adobe Firefly, Clay, Databricks, Gamma, Genspark, Glean, Krea, Listen Labs, Lovable, Suno, Synthesia) were all read on other pages of their first-party domains, with no logins: by text fetch, and in the cloud browser on public pages for Adobe Firefly, Suno, Glean and Listen Labs, where the fetched text had no prices. The aim was to see whether a credit price and allowance are public at all. Rate levels used throughout: R1, a dollar rate stated for the credit on the coded pages (9 of 22 products); R2, adding a rate derived from a plan price on the coded pages (ElevenLabs, Runway), 11 of 22; R3, adding a plan price with an allowance in the same credit found on other same-domain pages in this read (Adobe Firefly, Clay, Gamma, Genspark, Krea, Lovable, Suno, Synthesia), 19 of 22. Databricks is listed in Appendix B3 but not counted, because its only rate is a worked example for a different product page than the coded rows (counting it would make R3 20 of 22). Glean and Listen Labs showed no dollar

figure in the text read, which does not show that none exists. Listen Labs' own user guide says its platform fee is priced by outlined commitment, and its interview credits have no dollar rate in the text read; Glean's pricing page offers a demo and says nothing about how credits are priced, so Glean is not classed as contract-priced. Each R3 entry is in Appendix B3 with URL, date, quoted lines, billing route, allowance and hash. Unit match: each supplementary credit name is the unit of the coded rows, with the caveats in Appendix B3 (Clay data credits only, priced apart from actions; a Databricks DBU worked example on a different product page; plans that bundle features; and the Lovable and Synthesia notes on what a credit buys). R3 is rate availability. No supplementary price was paired with a coded action to form a one-event quote, no row was recoded, and the quote counts (5 of 22 and 6 of 22) are unchanged. Every product without a same-page rate was read; products with a same-page rate were not re-read, so effort is equal across the products where R3 could change the count.

Evidence and truncation. Each row stores a source URL, capture file with line span and the first 12 hex characters of the SHA-256 of the captured text (the full SHA-256 was also recorded). Page capture was limited to about 50 KB per page section. The full Snowflake PDF text (185 KB) is a complete retrieval, recorded as a full retrieval, so its size does not imply truncation. A row is flagged NEAR_CAP when its capture file is 49,000 bytes or more: 11 of 86 RQ1 rows and 1 of 25 RQ3 records. For these, absence is a lower bound and a quoted amount stays valid. 254 page sections were captured (40 on Nov 3 and 214 on Nov 4). Summary files are not pages and are excluded. A bounded recapture on Nov 4 added the Microsoft Cowork and Cost Management pages, the World Labs beta billing page and Snowflake pages, and the full Snowflake Service Consumption Table PDF was also added.

Reproduction and validation. The flags, rows, chains, frames and exclusion sensitivities and a per-product inventory were computed by script from the coded rows, with the primary population and the broadened frame kept apart. A validation script checks each coded row against the saved page text. It fails on a missing capture, a content hash that differs from the recorded hash, checked on the full SHA-256 (all 254 match), a captured page that is not recorded, a wrong sha12, an invalid or empty code, a line span beyond the page text, a missing field, a duplicate row key, or a worked-chain quote that is not verbatim on its cited line. It warns when numbers or unit words in the amount are not in the cited span. These checks cover integrity, references and schema. Duplicate keys are checked for RQ1 rows only, and schema is checked after the files open. The checks do not show that a code is right, and they make no semantic validation claim. Extraction is bounded to the features named on the pages read, so it is not a census of any product. For worked chains, the amount quote and the peg quote must appear on their cited lines, and the product of units and rate is compared with a page-stated dollar amount only where the page states one; the rest are arithmetic only, with conditions in Table 2. A separate check compares every Appendix B1 row in this document with the coded rows by line span and sha12: 114 of 114 rows match, and the one coded row absent from Appendix B1 is the removed Synthesia re-render row. Appendix B3 rows are checked only for quoted lines present in the saved page text, by the same quote-on-line rule. A reader can recount most

totals from Appendices B1, B2, B4 and B5 and can check any quote at its public URL and line span. The hash checks, the not-located results and the frame sensitivities in Table 3 depend on the saved page text and the frame rules, which are not part of the paper.

Manual source reread. On Nov 4, 2026 the single coder reread the cited text of every counted row: D and C rows against the action, amount and unit as coded, and P, N and I rows against whether the span supports a no-amount or incomplete status. Each reread is recorded on the row. The reread corrected rows where a cited line did not support the code as first recorded. This is one coder rereading its own work, not independent validation, and the capture returned partial pages for some sources.

Populations. The primary population is the original dataset row with the exact coded surface (22 products). Every headline result in Section 4 (RQ1 codes, precision, quantity and coverage, RQ2 chains, RQ3) uses this population. The broadened family frame (28 coded products, 35 product-surface pairs with RQ1 rows) adds three products that are different products from their dataset rows (Datadog Bits AI, Devin, Intercom Operator), two coded from a different surface (ChatGPT Business, Enterprise and Edu credit rate card; Copilot Credits for Copilot Studio and usage-based features) and Relevance AI. It is shown in a labeled secondary table. A product is a company or dataset row and a surface is the plan, interface or API a page describes, so one product can have several surfaces: RQ3 has 25 surface records for the 22 primary products. No coded surface is reported as a finding about a dataset row it does not match.

Precision rule. A single number that the page labels approximate, estimated or marks with a tilde is D with precision approximate (ElevenLabs, ChatGPT rate card, Datadog Bits). A range, a median and 90th percentile pair, a from-price, or a figure that varies by complexity is C. Model or size variants collapse into one row only when they share one precision class. A variant with a variable or from-price gets its own row. Worked examples are excluded when the underlying rates are coded.

Quantity field. yes: fixed per call, event or action. user_set: exact rate per unit the user sets (second, minute, character, megapixel, clip). no: depends on run-time use (tokens, neurons, GPU hours, mixed units, variable world size). Chain C1 needs yes. C1u needs user_set. Both need exact credits and a same-surface dollar rate.

Coder. One coder, working with AI assistance. No inter-rater reliability is claimed.

8 Limitations

- One coder applied the codes with AI assistance, and there is no independent human second coder and no independence claim. Codes are judgments.
- Snowflake Table 6 is coded as a selected, non-exhaustive subset (one model per function type and the named non-token features), so it is not a complete rate audit.
- Public pages only. Logged-in screens were not seen, so RQ3 levels understate what products do.
- Commercial terms not coded. The study codes only whether statements on expiry, rollover, running out of credits and failed runs were located (Appendices B4 and B5), not their effect on cost. It does not code true-up mechanics, seat plus credit

packaging, or currency and regional differences. These terms can change how much a documentation gap matters, and they are outside the scope of this audit. Buyer-side spend controls, meaning budget limits, spending caps, alerts and forecasting tools, and whether they work, were not evaluated, so the recommendations should not be read as implying that these controls were studied.

- Convenience cohort: 22 primary products (28 product-surface pairs) and 28 broadened coded products (35 pairs). No prevalence or ranking claims.
- Row counts depend on the row rule and page structure. The flag is an existence indicator and coverage is measured only over recorded rows.
- H1 is illustrated with two context pairs and is not tested. Plan prices bundle seats and features.
- The 75-row frame is the company list of an earlier study by the same author, and the saved version used was not recorded. The candidates and their status are listed, and every captured page is hashed, but no search log was kept for the page-finding step, so that step cannot be repeated exactly. The cohort is a defined convenience cohort.
- An earlier Snowflake Table 6 gap was resolved from the full official PDF (effective Oct 29, 2026, 185 KB, complete retrieval, SHA-256 prefix ef326e45e454). Some sources could not be read: the World Labs per-task rate card, which the public page says is shown only inside the platform. The Microsoft licensing guide capture is near the fetch cap, so its later sections were not read, but the Microsoft usage and Cost Management pages were.
- Fetch cap near 50 KB, flagged per row. 3 primary products and 4 broadened products have an incomplete-capture row.
- Pages change. Each row is a dated snapshot with a hash.
- No cross-product or cross-surface credit equivalence. The ChatGPT, Microsoft Copilot and World Labs surface mappings are stated in the scope trace and in Table 4. In the frozen baseline, plan prices on other pages of the same site were not followed, so a plan price shown elsewhere is not counted as a same-surface dollar rate. The dated Nov 6 supplementary read followed them for all 11 products without a same-page rate, and reports them separately in Appendix B3.
- 15 candidates had no credit pricing located, which reflects search limits.
- Pilot-informed freeze with full recoding after review.
- Validation checks references and arithmetic. The semantic reread was done by the same coder. No independent human audit is claimed.
- The literature review is bounded and not systematic. Novelty is not claimed.

9 Ethics and release

Public documentation only. No logins, accounts, paid tests or contact with the companies. Any company affiliated with the author is excluded from the coded products.

Use of AI tools. AI tools were used in this work to collect and capture the public sources, assist to code the rows, write the analysis scripts, edit and review the text

and check the work. The author set the research questions, scope and measures, and is responsible for the contents of this paper.

References

[1] Kaya, M. "They publish what they're selling, not what you're buying" (Pricing Transparency Audit). AI Tools Police, version 1.4, Sep 4, 2026. DOI 10.5281/zenodo.22303623 (all versions: 10.5281/zenodo.22234975). Accessed Nov 4, 2026. https://aitoolspolice.com/research/pricing-transparency-audit/

[2] Strategy of Finance. The AI Credit Index, Edition 1 (Q2 2026). Published Jun 13, 2026; observed Jun 2 to Jun 29, 2026. https://www.strategyoffinance.com/benchmarks/ai-credit-index/editions/edition-1/

[3] Strategy of Finance. How transparent is AI pricing (AI Credit Index, Edition 2 transparency page). Accessed Nov 6, 2026. https://www.strategyoffinance.com/benchmarks/ai-credit-index/transparency/

[4] UsagePricing. The credit equation is published in halves (trend page, marked superseded as of Sep 25, 2026). Accessed Nov 2026. https://www.usagepricing.com/blueprint/trends/credit-equation-published-in-halves

[5] TrueUpp Research. We read 64 usage-based pricing pages. Here's what they don't say. Oct 9, 2026. https://trueupp.com/research/pricing-terms-2026

[6] Sirvi, M. (IVRIS Tech). AI Pricing Transparency: 100-Vendor Source Audit. Version 1.1, Jul 26, 2026. https://ivristech.com/ai-pricing-transparency-b2b-software/

[7] DataDab. How B2B SaaS prices AI: credits beat add-ons, and premiums hide (The AI Pricing Premium Index, 106-company coding study). Oct 2026. https://www.datadab.com/research/b2b-saas-ai-pricing

[8] Gonzalez, J. D. Moving Targets: Promissory Instability in the Pricing and Access Terms of Generative AI Subscription Services (2024-2026). SSRN 7162638, written Jul 22, 2026, posted Aug 5, 2026. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=7162638

[9] PricingIO. Credit-Based AI Pricing: Why Buyers Can't Forecast It - and How to Fix That. Jul 15, 2026. https://www.pricingio.com/insights/credit-based-ai-pricing-why-buyers-cant-forecast-it-and-how-to-fix-that

[10] Sun, G., Wang, Z., Zhao, X., Tian, B., Shen, Z., He, Y., Xing, J., and Li, A. Position: Invisible Tokens, Visible Bills: The Urgent Need to Audit Hidden Operations in Opaque LLM Services. In Proceedings of the 43rd International Conference on Machine Learning, Proceedings of Machine Learning Research 306, pp. 172196-172207, 2026. https://proceedings.mlr.press/v306/sun26bd.html

[11] Yue, S., Cao, S., Wu, X., Li, R., Bai, Y., and Shi, J. Observable Correlates of Token Pricing in First-Party Large-Model API Services. In Proceedings of the 2026 3rd Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence, pp. 313-319, ACM, Sep 15, 2026. DOI 10.1145/3837873.3837924. https://dl.acm.org/doi/full/10.1145/3837873.3837924

[12] FinOps Foundation. Budgeting (FinOps Framework capability). Read Nov 6, 2026. https://www.finops.org/framework/capabilities/budgeting/

[13] FinOps Foundation. Forecasting (FinOps Framework capability). Read Nov 6, 2026. https://www.finops.org/framework/capabilities/forecasting/
[14] FinOps Foundation. Allocation (FinOps Framework capability). Read Nov 6, 2026. https://www.finops.org/framework/capabilities/allocation/
[15] FinOps Foundation. FinOps for AI Overview. Read Nov 6, 2026. https://www.finops.org/wg/finops-for-ai-overview/
[16] FinOps Foundation. How to Build a Generative AI Cost and Usage Tracker. Read Nov 6, 2026. https://www.finops.org/wg/how-to-build-a-generative-ai-cost-and-usage-tracker/
[17] Bala, R. and Carr, S. Usage-based pricing of software services under competition. Journal of Revenue and Pricing Management 9(3):204-216, 2010. DOI 10.1057/rpm.2010.12. https://doi.org/10.1057/rpm.2010.12
[18] Gupta, P. From Tokens to Credits to Outcomes: Pricing AI Across Models, Applications, and Cloud Platforms. arXiv:2610.30701 (v2), 2026. https://arxiv.org/abs/2610.30701

Appendix A. Scope trace for the 76 candidates

Candidate (dataset row)	**Status**	**Basis**	**Mapping note**
OpenAI (API)	WALLET (listed, not coded)	balance is dollars, single currency	
OpenAI (ChatGPT consumer)	IN (credit meter)	ChatGPT credit-based plans and rate card; page-read	Dataset row is consumer ChatGPT. Pages coded describe the Business/ Enterprise/ Edu credit-based rate card. Surface mismatch: coded surface is not the consumer plans.
Anthropic (Claude API)	WALLET (listed, not coded)	balance is dollars, single currency	
Anthropic (Claude consumer)	WALLET (search excerpt)	usage credits are pay-as-you-go balance after plan limits, bundles discounted versus standard usage credit rates (https://support.claude.com/en/articles/12429409-manage-usage-credits-for-paid-claude-plans, https://support.claude.com/en/articles/14246112-buy-usage-bundles); evidence is search excerpt only, article not read in full	
Cursor (Anysphere)	OUT (page-verified)	included usage is an unnamed weekly pool with relative multiples (4x, 15x, 50x Plus); extra usage is dollar prepaid credit at API list price (https://cursor.com/docs/models-and-pricing); boundary case	
Mistral AI	WALLET (listed, not coded)	prepaid credits are an amount in billing currency (https://docs.mistral.ai/admin/billing-usage/billing); tokens priced per million	

Candidate (dataset row)	**Status**	**Basis**	**Mapping note**
Cohere	OUT (page-verified)	no credit unit in captured pricing pages (https://cohere.com/pricing)	
Reflection AI	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Safe Superintelligence (SSI)	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Thinking Machines Lab	OUT (page-verified)	per million tokens and per GB-month (https://thinkingmachines.ai/tinker/pricing/)	
Databricks	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Crusoe	OUT (page-verified)	per GPU-hour prices (https://www.crusoe.ai/cloud/pricing)	
SambaNova	CAPTURE INCOMPLETE	pricing page 410 bytes (https://cloud.sambanova.ai/plans/pricing)	
Together AI	WALLET (page-verified)	prepaid credits are a dollar balance (minimum $5 purchase) (https://support.together.ai/articles/1862638756-changes-to-free-tier-and-billing-july-2025)	
Baseten	WALLET (page-verified)	credits applied to per-token and per-minute usage (https://docs.baseten.co/organization/billing)	
Fireworks AI	WALLET (page-verified)	usage draws down prepaid credits at per-token and GPU prices (https://fireworks.ai/pricing)	
Fal	WALLET (page-verified)	model APIs priced per output drawn from prepaid credits; one model lists '$0.01 per credits' (https://fal.ai/pricing); boundary case	
Cognition (Devin + Windsurf)	IN (credit meter)	ACUs and prepaid on-demand credits; page-read	Devin pages only; Windsurf not coded.
Replit	WALLET (page-verified)	pricing page: '$20 towards more powerful models' on Core, i.e. a dollar credit; credit packs bought as dollars (https://replit.com/pricing, https://docs.replit.com/billing/managing-spend); Replit Cells rows kept as record, not counted	
Abridge	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Chai Discovery	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
OpenEvidence	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Physical Intelligence	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Skild AI	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Applied Intuition	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Perplexity	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Glean	IN (credit meter)	FlexCredits at a rate card; page-read	
Notion	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Genspark	IN (credit meter)	named credits; page-read	

Candidate (dataset row)	Status	Basis	Mapping note
Gamma	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Listen Labs	IN (credit meter)	Interview Credits per completed interview; page-read	
Speak	OUT (page-verified)	subscription; no credit unit in captured pages (https://www.speak.com/)	
Harvey	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Legora	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Sierra	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://sierra.ai/blog/outcome-based-pricing-for-ai-agents	
Decagon	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
EliseAI	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Rogo	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Clay	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Midjourney	OUT (page-verified)	fast GPU time in hours, extra GPU time $4/hr (https://docs.midjourney.com/hc/en-us/articles/27870484040333-Comparing-Midjourney-Plans); resource-time metering, not a named credit unit; boundary case	
ElevenLabs	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Runway	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Synthesia	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
HeyGen	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Black Forest Labs	IN (credit meter)	dollar-pegged credit with per-generation burn; page-read	
Krea	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Suno	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Cyera	CAPTURE INCOMPLETE	pricing page 866 bytes (https://www.cyera.com/pricing)	
Lovable	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
World Labs	IN (credit meter)	credits per generation; page-read	
Mercor	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	
Surge AI	NO CREDIT PRICING LOCATED	bounded first-party search (search plus up to 3 pages) found no credit or unit pricing; this is a search limit, not a finding about the product	

Candidate (dataset row)	**Status**	**Basis**	**Mapping note**
Snowflake (AI credits)	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	Cortex AI credit rows coded from the full Service Consumption Table (Table 2(b) and Table 6); two Platform Credit rows kept but not counted.
AWS Bedrock	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://aws.amazon.com/bedrock/pricing/ (fetch truncated at 50 KB)	
Google Vertex AI	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://cloud.google.com/model-garden/pricing (fetch truncated at 50 KB)	
Azure AI (incl. Azure OpenAI)	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://azure.microsoft.com/en-us/pricing/details/cognitive-services/openai-service	
Cloudflare Workers AI	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Vercel (AI Cloud)	WALLET (listed, not coded)	monthly dollar usage credit applied to metered dollar prices	
MongoDB Atlas (Vector Search)	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://www.mongodb.com/pricing (credits mentioned are startup credits)	
Pinecone	WALLET (page-verified)	prepaid credits are dollars ($8,000 to $25,000) against read, write and storage units (https://docs.pinecone.io/guides/manage-cost/understanding-cost)	
Supabase (pgvector/AI)	WALLET (listed, not coded)	dollar usage credit applied to dollar prices	
Confluent Cloud (AI)	WALLET (page-verified)	$400 signup credit applied to dollar usage (https://docs.confluent.io/cloud/current/billing/billing-faq.html)	
Datadog (LLM Observability)	IN (credit meter)	AI Credits; page-read	Datadog AI Credits (Bits AI) pages coded; LLM Observability SKU not coded.
CoreWeave	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://www.coreweave.com/pricing	
Lambda	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://lambda.ai/pricing	
Nebius	OUT (page-verified)	per-token, per-hour or outcome pricing; no source credit unit in captured pages: https://nebius.com/prices (credit mentions are promo credits)	
GitHub Copilot	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
Salesforce Agentforce	IN (credit meter)	Flex Credits per action; page-read	

Candidate (dataset row)	Status	Basis	Mapping note
Intercom Fin	IN (credit meter)	credits for Operator plus per-outcome Fin; page-read	Operator credits page coded; Fin is priced per resolution (outcome), not credits.
Zendesk AI	OUT (page-verified)	outcome-based unit (automated resolution tiers, resolution allowance), not a credit unit (https://support.zendesk.com/hc/en-us/articles/9570369117338-About-automated-resolution-tiers); outcome pricing is a scope boundary	
Microsoft 365 Copilot	IN (credit meter)	Copilot Credits (Copilot Studio); page-read. Row is the dataset's per-seat M365 Copilot; credits apply to Copilot Studio and usage-based features	Copilot Credits licensing guide coded. Copilot Studio rows are a different surface and are excluded from product counts.
Adobe Firefly	IN (credit meter)	named non-currency unit or credit with a stated unit-to-dollar link	
DeepSeek	OUT (page-verified)	per-million-token prices (https://api-docs.deepseek.com/quick_start/pricing/)	
Replicate	WALLET (page-verified)	prepaid credit applied to per-output and per-second prices (https://replicate.com/docs/topics/billing/prepaid-credit/)	
OpenRouter	WALLET (page-verified)	credits are deposits deducted at provider list price (https://openrouter.ai/docs/faq)	
Relevance AI (extension, not in dataset)	IN (credit meter)	Actions plus Vendor Credits (product term) as separate meters	

Appendix B1. RQ1 action-family rows

Each line: [code/precision/quantity known] surface | action family | amount as stated | same-surface bridge | line span in the captured page | sha12 of the capture, then the source link.

Runway

- [D/exact/qty user_set] Runway web app credits | Video generation by model (Gen-4.5, Aleph 2.0, Kling 4.0, Seedance 2.0 Pro) | credits per clip by model, stated per clip duration (Gen-4.5 60 per 5 s; Aleph 2.0 140 per 5 s; Kling 4.0 118 per 5 s; Seedance 2.0 Pro 160 per 4 s) | bridge: credit-to-$ not stated as a rate on this page; plan prices vs credits shown (lines 26-44,68) | lines 54-74 | sha12 2482b375fe91 Source: https://runway.com/pricing

ElevenLabs

- [D/approximate/qty user_set] ElevenLabs web app credits | Text to speech, V2 Multilingual | 1 credit per character (page says approximate costs) | bridge: plan table: Starter $6 for 30,000 credits etc (line 202) | lines 206 | sha12 6fe96760808d Source: https://elevenlabs.io/pricing
- [C/range_or_estimate/qty user_set] ElevenLabs web app credits | Text to speech, V2 Flash/Turbo and V2.5 Flash/Turbo (API) | 0.5 to 1 credit per character (page says approximate costs) | bridge: plan table: Starter $6 for 30,000 credits etc (line 202) | lines 206 | sha12 6fe96760808d Source: https://elevenlabs.io/pricing
- [D/approximate/qty user_set] ElevenLabs web app credits | Speech to text | 330 credits per minute (page says approximate costs) | bridge: plan table: Starter $6 for 30,000 credits etc (line 202) | lines 210 | sha12 6fe96760808d Source: https://elevenlabs.io/pricing
- [D/approximate/qty user_set] ElevenLabs web app credits | Eleven Music | 900 credits per minute (page says approximate costs) | bridge: plan table: Starter $6 for 30,000 credits etc (line 202) | lines 210 | sha12 6fe96760808d Source: https://elevenlabs.io/pricing
- [D/approximate/qty yes] ElevenLabs web app credits | Sound effects | 200 credits per generation (page says approximate costs) | bridge: plan table: Starter $6 for 30,000 credits etc (line 202) | lines 210 | sha12 6fe96760808d Source: https://elevenlabs.io/pricing
- [D/approximate/qty user_set] ElevenLabs web app credits | Voice changer and voice isolator | 1,000 credits per minute (page says approximate costs) | bridge: plan table: Starter $6 for 30,000 credits etc (line 202) | lines 210 | sha12 6fe96760808d Source: https://elevenlabs.io/pricing
- [D/approximate/qty user_set] ElevenLabs web app credits | Dubbing (automatic with watermark / without / Studio with / Studio without) | 2,000 / 3,000 / 5,000 / 10,000 credits per minute (page says approximate costs) | bridge: plan table: Starter $6 for 30,000 credits etc (line 202) | lines 210 | sha12 6fe96760808d Source: https://elevenlabs.io/pricing

Synthesia

- [D/exact/qty user_set] Synthesia self-serve credits | Video generation | 2 credits per second (one minute = 120) | bridge: plan credit allowances only; no credit-to-$ rate on this page | lines 16 | sha12 2ae39dfc9bc7 Source: https://help.synthesia.io/en/articles/8850453-how-does-synthesia-credit-consumption-work-for-video-edits

HeyGen

- [D/exact/qty user_set] HeyGen credit-based plans | Avatar video generation by avatar type (III, IV, V, custom motion) | 4 to 48 credits per minute by avatar type | bridge: Business extra credits $0.05 per credit (line 41) | lines 52-57 | sha12 6dd246da1901 Source: https://help.heygen.com/en/articles/15126059-how-to-use-credits-on-heygen

- [D/exact/qty yes] HeyGen credit-based plans | AI model training | 60 credits per model | bridge: Business extra credits $0.05 per credit (line 41) | lines 65 | sha12 6dd246da1901 Source: https://help.heygen.com/en/articles/15126059-how-to-use-credits-on-heygen
- [D/exact/qty user_set] HeyGen credit-based plans | Video translation, speed mode / precision mode | 6 / 10 credits per minute | bridge: Business extra credits $0.05 per credit (line 41) | lines 68 | sha12 6dd246da1901 Source: https://help.heygen.com/en/articles/15126059-how-to-use-credits-on-heygen

Perplexity

- [C/range_or_estimate/qty no] Perplexity Computer credits (consumer plans) | Quick fact or definition | 15 to 70 credits (page: examples, not guarantees) | bridge: 100 credits = $1 (line 13) | lines 13,17-50 | sha12 d88eb4abdb35 Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity
- [C/range_or_estimate/qty no] Perplexity Computer credits (consumer plans) | Multi-step research report | 350 to 950 credits (page: examples, not guarantees) | bridge: 100 credits = $1 (line 13) | lines 13,17-50 | sha12 d88eb4abdb35 Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity
- [C/range_or_estimate/qty no] Perplexity Computer credits (consumer plans) | Data dashboard | 2,400 to 9,800 credits (page: examples, not guarantees) | bridge: 100 credits = $1 (line 13) | lines 13,17-50 | sha12 d88eb4abdb35 Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity
- [C/range_or_estimate/qty no] Perplexity Computer credits (consumer plans) | Entire app or multi-day build | 15,000+ credits (page: examples, not guarantees) | bridge: 100 credits = $1 (line 13) | lines 13,17-50 | sha12 d88eb4abdb35 Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity

Notion

- [C/range_or_estimate/qty no] Notion credits (Agent, Standard and Select allowances) | Simple task (Auto / Opus 5.5) | 5 / 50 estimated credits per task; use varies | bridge: $10 per 1,000 credits monthly, $13 annual (annual billing is the higher rate) (lines 34-36) | lines 181-186 | sha12 199673d365b6 Source: https://www.notion.com/help/what-are-notion-credits
- [C/range_or_estimate/qty no] Notion credits (Agent, Standard and Select allowances) | Medium task (Auto / Opus 5.5) | 10 / 85 estimated credits per task; use varies | bridge: $10 per 1,000 credits monthly, $13 annual (annual billing is the higher rate) (lines 34-36) | lines 181-186 | sha12 199673d365b6 Source: https://www.notion.com/help/what-are-notion-credits
- [C/range_or_estimate/qty no] Notion credits (Agent, Standard and Select allowances) | Complex task (Auto / Opus 5.5) | 25 / 150 estimated credits per

task; use varies | bridge: $10 per 1,000 credits monthly, $13 annual (annual billing is the higher rate) (lines 34-36) | lines 181-186 | sha12 199673d365b6 Source: https://www.notion.com/help/what-are-notion-credits

Clay

- [D/exact/qty yes] Clay Data Credits | Find email (provider-dependent) | 2 credits (Prospeo, DropContact, Datagma, Hunter.io); 5 credits (People Data Labs) | bridge: plan credits only on this page | lines 15 | sha12 347c1e000f2d Source: https://www.clay.com/faq/what-are-clay-credits-and-how-do-they-work
- [D/exact/qty yes] Clay Data Credits | Find mobile number | starts at 2; 5 (PDL), 13 (Selligence), 15 (ContactOut), 25 (Datagma) | bridge: none stated | lines 17 | sha12 347c1e000f2d Source: https://www.clay.com/faq/what-are-clay-credits-and-how-do-they-work
- [C/range_or_estimate/qty no] Clay Data Credits | Most other actions | free or 1 to 2 credits | bridge: Actions meter separate from Data Credits (line 610) | lines 13 | sha12 347c1e000f2d Source: https://www.clay.com/faq/what-are-clay-credits-and-how-do-they-work

Lovable

- [D/exact/qty yes] Lovable credits | Plan Mode message | 1 credit per message | bridge: no credit-to-$ rate located yet | lines 28 | sha12 8d7903152277 Source: https://lovable.dev/pricing
- [C/range_or_estimate/qty no] Lovable credits | Default Mode prompts (four worked examples) | 0.50, 0.90, 1.20 and 1.70 credits for four example prompts (varies with complexity) | bridge: none stated | lines 36-39 | sha12 8d7903152277 Source: https://lovable.dev/pricing

Relevance AI (dataset row: not in the dataset, added, not in dataset)

- [D/exact/qty yes] Relevance AI Actions | Tool run (one Action) | one Action per Tool run; failed runs still count | bridge: $80 per 1,000 Actions (line 33) | lines 13-14 | sha12 b9a2c97618b4 Source: https://relevanceai.com/docs/admin/subscriptions/plans
- [P/none/qty no] Relevance AI Vendor Credits (product term) | LLM and tool cost | exact cost, no markup; amounts depend on model and tokens; no per-action numbers on page | bridge: $20 per 10,000 Vendor Credits (product term) (line 34) | lines 19-21 | sha12 b9a2c97618b4 Source: https://relevanceai.com/docs/admin/subscriptions/plans

Adobe Firefly

- [D/exact/qty yes] Adobe Firefly generative credits (Firefly FAQ page) | Most standard generative features (e.g. Generative Fill) | 1 credit per generation; some standard features exceed 1 (table elsewhere) | bridge: plan credit allowances only

| lines 699 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html

- [D/exact/qty user_set] Adobe Firefly generative credits (Firefly FAQ page) | Generate Video 1080p/720p/540p at 24 FPS | 100 / 50 / 20 credits per second | bridge: none stated | lines 743-751 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [D/exact/qty user_set] Adobe Firefly generative credits (Firefly FAQ page) | Generative Extend in Premiere (4K 30 175, 1080p 30 125, 720p 30 75, 4K 24 150, 1080p 24 100, 720p 24 50) | credits per second as listed | bridge: none stated | lines 757-777 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [D/exact/qty user_set] Adobe Firefly generative credits (Firefly FAQ page) | Translate Video / Translate Audio | 5 credits per second (4K 15; Translate & Lip Sync 1080p 10) | bridge: none stated | lines 789-799 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [D/exact/qty yes] Adobe Firefly generative credits (Firefly FAQ page) | Generate Sound Effects | 10 credits per generation | bridge: none stated | lines 801-805 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [D/exact/qty user_set] Adobe Firefly generative credits (Firefly FAQ page) | Text to Avatar (beta) | 10 credits per second | bridge: none stated | lines 807-811 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [D/exact/qty user_set] Adobe Firefly generative credits (Firefly FAQ page) | Generate Music | 20 credits per minute | bridge: none stated | lines 813-815 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [D/exact/qty yes] Adobe Firefly generative credits (Firefly FAQ page) | Firefly Graph: mask, refine mask, RGBX nodes | 10 credits per image | bridge: none stated | lines 888-910 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [D/exact/qty yes] Adobe Firefly generative credits (Firefly FAQ page) | Firefly Graph: GPT 5.1 and Gemini 2.5 LLM nodes | 0 credits per image for a limited time | bridge: none stated | lines 914-921 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [I/none/qty n/a] Adobe Firefly generative credits (Firefly FAQ page) | Partner models (Google, OpenAI, ElevenLabs, Topaz and others) | charged per separate partner-models page, not captured | bridge: none stated | lines 737 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html
- [N/none/qty n/a] Adobe Firefly generative credits (Firefly FAQ page) | Firefly Graph workflows containing LLM nodes or loops (undetermined-cost nodes) | Graph shows "Credits vary"; cost cannot be known until the workflow finishes | bridge:

none stated | lines 930-949 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html

Glean

- [D/exact/qty yes] Glean Enterprise Flex (FlexCredits) | Basic Search Query | 1 (Client APIs); 0 unlimited from Assistant or Embedded Apps FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Fast Mode Assistant Query | variable; typical ˜3 (p50), ˜15 (p90) FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Thinking Mode Assistant Query (basic/standard and premium models) | basic/standard ˜7 (p50), ˜26 (p90); premium ˜35 (p50), ˜120 (p90) FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Adaptive Reasoning Query (basic/standard and premium) | basic/standard ˜11 / ˜38; premium ˜26 / ˜83 FlexCredits (p50 / p90) | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Voice Session | ˜3 / ˜26 FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Meeting Notes (per meeting) | ˜9 / ˜18 FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Code Writer Query | ˜9 / ˜32 FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Image Generation Query | ˜7 / ˜9 FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Slide Generation Query | ˜45 / ˜142 FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Memory and personalization (per monthly active user) | ˜5 / ˜10 FlexCredits | bridge: no public

dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing
- [C/range_or_estimate/qty no] Glean Enterprise Flex (FlexCredits) | Glean Agent Run | ~7 / ~114 FlexCredits | bridge: no public dollar price per FlexCredit located on this page | lines 85-101 | sha12 3bb1d1360f26 Source: https://docs.glean.com/glean-enterprise-flex-pricing

Gamma

- [C/range_or_estimate/qty yes] Gamma API (developers.gamma.app) | Text generation per card | 1 to 3 credits per card | bridge: API and app draw from the same credit balance (line 55); no credit-to-$ rate on this page | lines 70 | sha12 3f25d6aee601 Source: https://developers.gamma.app/get-started/access-and-pricing
- [C/range_or_estimate/qty yes] Gamma API (developers.gamma.app) | AI image generation by model tier (API) | 2 to 125 credits per image by model tier | bridge: API and app draw from the same credit balance (line 55); no credit-to-$ rate on this page | lines 71 | sha12 3f25d6aee601 Source: https://developers.gamma.app/get-started/access-and-pricing
- [C/range_or_estimate/qty yes] Gamma API (developers.gamma.app) | Worked examples (deck 10 cards + 5 images; doc 20 cards + 15 images; socials 30 + 30) | ~20-60; ~320-1,070; ~930-3,900 credits | worked example built from the per-card and per-image rates already coded; excluded to avoid double counting | bridge: API and app draw from the same credit balance (line 55); no credit-to-$ rate on this page | lines 79-81 | sha12 3f25d6aee601 | NOT COUNTED Source: https://developers.gamma.app/get-started/access-and-pricing
- [P/none/qty n/a] Gamma API (developers.gamma.app) | Agent generations and edits (quality presets) | page: higher presets cost more; no amounts | bridge: none stated | lines 95 | sha12 3f25d6aee601 Source: https://developers.gamma.app/get-started/access-and-pricing
- [P/none/qty n/a] Gamma app (help center) | Image generation and editing in the app | help page: the credit cost for the model is shown before an image is generated or edited; no amounts on pages read | bridge: none stated | lines 22 | sha12 4c4ff9d20ddb Source: https://help.gamma.app/en/articles/7834324-how-do-credits-work-in-gamma

GitHub Copilot

- [D/exact/qty no] GitHub Copilot AI Credits (usage-based, over included allowance) | Per-model token rates (input, cached input, output) in pricing tables | $ per 1M tokens by model (e.g. promotional Gemini 3.8 Flash $0.75 input, $3.75 output); credits = dollars / $0.01 | bridge: same page: 1 AI credit = $0.01 (lines 11,20) | lines 11-213 | sha12 16d48da537e6 Source: https://docs.github.com/en/copilot/reference/copilot-billing/models-and-pricing
- [P/none/qty no] GitHub Copilot AI Credits (usage-based, over included allowance) | Credits for a whole interaction or task | page: depends on model and tokens; no per-task credit figure | bridge: same page: 1 AI credit = $0.01 (lines 11,20) | lines 11 |

sha12 16d48da537e6 Source: https://docs.github.com/en/copilot/reference/copilot-billing/models-and-pricing

Suno

- [D/exact/qty yes] Suno credits (help center) | Song generation (two songs) | 10 credits; more if many image or video inputs | bridge: no credit-to-$ rate on this page | lines 59 | sha12 e753a13ff860 Source: https://help.suno.com/en/articles/13924481
- [P/none/qty n/a] Suno credits (help center) | Generation with many image/video inputs | page says cost increases; no amount | bridge: none stated | lines 59 | sha12 e753a13ff860 Source: https://help.suno.com/en/articles/13924481
- [P/none/qty n/a] Suno credits (help center) | Max Mode | page says costs more credits; no amount | bridge: none stated | lines 70 | sha12 9ad1fc70ac0f Source: https://help.suno.com/en/articles/13924481

Cloudflare Workers AI

- [D/exact/qty no] Cloudflare Workers AI (Neurons) | Text models (per M input / output tokens), e.g. llama-3.1-8b-instruct-fp8-fast | 4,119 Neurons per M input, 34,868 per M output ($0.045 / $0.384 per M) | bridge: same page: $0.011 per 1,000 Neurons; 10,000 Neurons per day free (lines 14-25) | lines 49-54 | sha12 a19bee9c0bee Source: https://developers.cloudflare.com/workers-ai/platform/pricing/
- [D/exact/qty no] Cloudflare Workers AI (Neurons) | Image and audio models | per-unit rates differ by model and unit, e.g. flux-1-schnell 4.80 neurons per 512x512 tile and 9.60 per step; whisper 41.14 neurons per audio minute; flux-2-klein-9b priced per megapixel | bridge: same page: $0.011 per 1,000 Neurons; 10,000 Neurons per day free (lines 14-25) | lines 104-109,115-121 | sha12 a19bee9c0bee Source: https://developers.cloudflare.com/workers-ai/platform/pricing/

Snowflake (AI credits)

- [D/exact/qty no] Snowflake Service Consumption Table (AI Credits and Platform Credits) | Platform compute by warehouse size (Platform Credits, not AI Credits; out of AI-credit scope) | Platform Credits per hour by cloud and warehouse size (lines 86-108) | bridge: none stated | lines 86-108 | sha12 ab05db03b11b | NEAR_CAP | NOT COUNTED Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf
- [D/exact/qty yes] Snowflake Service Consumption Table (AI Credits and Platform Credits) | Cloud services compute (Platform Credits, not AI Credits; out of AI-credit scope) | 4.4 credits per compute-hour (line 32), page near cap | bridge: none stated | lines 32 | sha12 ab05db03b11b | NEAR_CAP | NOT COUNTED Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf

Krea

- [D/exact/qty yes] Krea image generation (docs, compute units) | Image models table: Krea 1, Krea 2 Turbo, Flux, Nano Banana, Seedream 3/4, Wan 2.5, Qwen,

Imagen 3/4/4 Fast/4 Ultra, Runway Gen-4, Ideogram 3.0, Flux Kontext/Pro, Flux 1.1 Pro/Ultra (18 models) | compute units per generation, e.g. Krea 1 8, Krea 2 Turbo 2, Flux 5, Seedream 4 24, Nano Banana 114, Wan 2.5 569 | bridge: plan allowances in units (5,000 / 20,000 / 60,000 per month); no per-unit dollar rate located on these pages | lines 47-66 | sha12 cc2101d86ae3 Source: https://www.krea.ai/docs/user-guide/features/krea-image

- [P/none/qty n/a] Krea image generation (docs, compute units) | Krea 3 | table says "See model picker" for compute units | bridge: none stated | lines 48 | sha12 cc2101d86ae3 Source: https://www.krea.ai/docs/user-guide/features/krea-image
- [I/none/qty n/a] Krea other products (video, enhance, edit) | Non-image generations | Video and Enhance tools are named; compute units for them are not on the page read | bridge: none stated | lines 214-215 | sha12 cc2101d86ae3 Source: https://www.krea.ai/docs/user-guide/features/krea-image

World Labs

- [D/exact/qty yes] World Labs Marble (docs.worldlabs.ai/marble/models) | Marble 1.1 world generation | 1,500 credits per generation | bridge: no credit-to-$ bridge asserted: Marble and World API credits are not interchangeable | lines 18 | sha12 61a6b9ee7206 Source: https://docs.worldlabs.ai/marble/models
- [C/range_or_estimate/qty no] World Labs Marble (docs.worldlabs.ai/marble/models) | Marble 1.1 Plus world generation | 1,500 base plus 0 to 1,500 variable credits | bridge: as above | lines 14 | sha12 61a6b9ee7206 Source: https://docs.worldlabs.ai/marble/models
- [D/exact/qty yes] World Labs Marble (docs.worldlabs.ai/marble/models) | Fastest model generation | 150 credits per generation | bridge: as above | lines 22 | sha12 61a6b9ee7206 Source: https://docs.worldlabs.ai/marble/models

Genspark

- [P/none/qty n/a] Genspark credits guide (help center) | All task types (chat, code, image, video, audio, agent) | cost drivers listed (model, length, resolution, duration, retries, auto-select fan-out); no credit amounts; "check the live pricing inside Genspark" | bridge: none stated | lines 62-146 | sha12 e040774aff42 Source: https://www.genspark.ai/helpcenter/credits-guide

Salesforce Agentforce

- [D/exact/qty yes] Salesforce Agentforce Flex Credits (pricing page) | Agentforce Action | 20 Flex Credits per action | bridge: same page: $500 per 100,000 Flex Credits (US) = $0.005 per credit (lines 34,172) | lines 156,172-192 | sha12 df9e9b9e067e Source: https://www.salesforce.com/agentforce/pricing/
- [C/range_or_estimate/qty yes] Salesforce Agentforce Flex Credits (pricing page) | Agentforce Voice use case (worked example) | worked example: 4 actions = 120 Flex Credits = US$0.15 per use case (30 per action derived); page says usage types and multipliers differ for Voice and points to the rate card | worked example; excluded (the per-action rate is coded separately) | bridge: same page: $500 per 100,000

Flex Credits (US) = $0.005 per credit (lines 34,172) | lines 245-257,295-297 | sha12 df9e9b9e067e | NOT COUNTED Source: https://www.salesforce.com/agentforce/pricing/

Microsoft Copilot Credits for Copilot Studio and usage-based features (dataset row: Microsoft 365 Copilot, surface differs from dataset row)

- [D/exact/qty yes] Microsoft Copilot Credits (licensing guide; dataset row "Microsoft 365 Copilot") | Work IQ Tools API call | 0.1 Copilot Credit per API call | bridge: same page: pay-as-you-go $0.01 per Copilot Credit (line 90) | lines 347 | sha12 1e5fe93212e7 | NEAR_CAP Source: https://www.microsoft.com/licensing/guidance/CopilotCreditsGuide
- [D/exact/qty yes] Microsoft Copilot Credits (licensing guide; dataset row "Microsoft 365 Copilot") | Copilot Managed Runtime hosting | 0.1 Copilot Credits per API call | bridge: same page: pay-as-you-go $0.01 per Copilot Credit (line 90) | lines 548-554 | sha12 1e5fe93212e7 | NEAR_CAP Source: https://www.microsoft.com/licensing/guidance/CopilotCreditsGuide
- [P/none/qty n/a] Microsoft Copilot Credits (licensing guide; dataset row "Microsoft 365 Copilot") | Work IQ Chat or Context API query | variable; driven by models, runtime, context and tools; no amounts | bridge: same page: pay-as-you-go $0.01 per Copilot Credit (line 90) | lines 346 | sha12 1e5fe93212e7 | NEAR_CAP Source: https://www.microsoft.com/licensing/guidance/CopilotCreditsGuide
- [C/range_or_estimate/qty no] Microsoft Copilot Credits (licensing guide; dataset row "Microsoft 365 Copilot") | GitHub Copilot harness natural-language creation (light 1-2 turns, medium 3-5, heavy 6+) | 1-20 / 21-60 / 61+ estimated credits | bridge: same page: pay-as-you-go $0.01 per Copilot Credit (line 90) | lines 400-408 | sha12 1e5fe93212e7 | NEAR_CAP Source: https://www.microsoft.com/licensing/guidance/CopilotCreditsGuide
- [I/none/qty n/a] Microsoft Copilot Credits (licensing guide; dataset row "Microsoft 365 Copilot") | Copilot Cowork prompts | estimation method stated (users x prompts x average price per prompt); the price table is an image, not captured as text | bridge: same page: pay-as-you-go $0.01 per Copilot Credit (line 90) | lines 203-209 | sha12 1e5fe93212e7 | NEAR_CAP Source: https://www.microsoft.com/licensing/guidance/CopilotCreditsGuide

(adjacent surface, not a dataset row, excluded from counts)

- [P/none/qty n/a] Microsoft Copilot Studio (Azure pay-as-you-go pricing page) | Agent responses and actions | credits decremented by task complexity, agent type and knowledge source; the price table renders "$-" in capture; no amounts | bridge: price per credit not rendered on this page; $0.01 appears on the licensing guide above (different surface page, same meter name) | lines 20-30 | sha12 7e5d2ca83f89 | NOT COUNTED Source: https://azure.microsoft.com/en-us/pricing/details/copilot-studio/

ChatGPT Business, Enterprise and Edu credit rate card (OpenAI) (dataset row: OpenAI (ChatGPT consumer), surface differs from dataset row)

- [D/approximate/qty yes] ChatGPT credit-based workspaces: Business/Enterprise/Edu rate card (help article); NOT the consumer plans | Chat message by model (10, 20 or 50 credits per message) | 10, 20 or 50 credits per message by model (page: approximately) | bridge: no credit-to-dollar rate on this page (credits are "billing units"; refer to invoice); bridge not established | lines 40-57 | sha12 bde1458b014e Source: https://help.openai.com/en/articles/11481834-chatgpt-rate-card-business-enterpriseedu-credit-based-pricing
- [D/approximate/qty yes] ChatGPT credit-based workspaces: Business/Enterprise/Edu rate card (help article); NOT the consumer plans | Agent mode | 30 credits per message (page: approximately) | bridge: no credit-to-dollar rate on this page (credits are "billing units"; refer to invoice); bridge not established | lines 65 | sha12 bde1458b014e Source: https://help.openai.com/en/articles/11481834-chatgpt-rate-card-business-enterpriseedu-credit-based-pricing
- [D/approximate/qty yes] ChatGPT credit-based workspaces: Business/Enterprise/Edu rate card (help article); NOT the consumer plans | Deep research | 50 credits per task (page: approximately) | bridge: no credit-to-dollar rate on this page (credits are "billing units"; refer to invoice); bridge not established | lines 66 | sha12 bde1458b014e Source: https://help.openai.com/en/articles/11481834-chatgpt-rate-card-business-enterpriseedu-credit-based-pricing
- [D/approximate/qty yes] ChatGPT credit-based workspaces: Business/Enterprise/Edu rate card (help article); NOT the consumer plans | Images | 5 credits per generation (page: approximately) | bridge: no credit-to-dollar rate on this page (credits are "billing units"; refer to invoice); bridge not established | lines 67 | sha12 bde1458b014e Source: https://help.openai.com/en/articles/11481834-chatgpt-rate-card-business-enterpriseedu-credit-based-pricing
- [D/approximate/qty user_set] ChatGPT credit-based workspaces: Business/Enterprise/Edu rate card (help article); NOT the consumer plans | Voice in ChatGPT | 1.25 credits per minute (page: approximately) | bridge: no credit-to-dollar rate on this page (credits are "billing units"; refer to invoice); bridge not established | lines 68 | sha12 bde1458b014e Source: https://help.openai.com/en/articles/11481834-chatgpt-rate-card-business-enterpriseedu-credit-based-pricing
- [D/approximate/qty no] ChatGPT credit-based workspaces: Business/Enterprise/Edu rate card (help article); NOT the consumer plans | Work and Codex usage by model (e.g. GPT-6.1 Sol: 50 input / 2.5 cached / 250 output per 1M tokens) | credits per 1M tokens; total depends on tokens (unknown until run) (page: approximately) | bridge: no credit-to-dollar rate on this page (credits are "billing units"; refer to invoice); bridge not established | lines 97-100,172-175 | sha12 bde1458b014e Source: https://help.openai.com/en/articles/11481834-chatgpt-rate-card-business-enterpriseedu-credit-based-pricing

Black Forest Labs

- [D/exact/qty yes] Black Forest Labs API and Playground credits | FLUX.1 family image generation and editing (fixed credits per image: Kontext, 1.1, Fill) | FLUX.1 Kontext [pro] 4 credits per image, [max] 8; FLUX1.1 pro 4, Ultra 6; Fill 5 (1 credit = $0.01) | bridge: same page: 1 credit = $0.01 (line 12) | lines 12,43-62 | sha12 9906a21a7c93 Source: https://help.bfl.ai/articles/7986977817-what-are-the-costs-associated-with-using-your-models
- [D/exact/qty user_set] Black Forest Labs API and Playground credits | FLUX 3 video per second (t2v, i2v full render hd/fhd; draft) | Full render $0.17/s hd or $0.29/s fhd (17 or 29 credits per second at 1 credit = $0.01); draft $0.06/s (6 credits); video continuation $0.43 to $0.54/s full, $0.12/s draft | bridge: same page: 1 credit = $0.01 (line 12) | lines 12,16-24 | sha12 9906a21a7c93 Source: https://help.bfl.ai/articles/7986977817-what-are-the-costs-associated-with-using-your-models
- [C/range_or_estimate/qty user_set] Black Forest Labs API and Playground credits | FLUX.2 image generation (priced per megapixel, minimum or from-prices, tiered rounding) | priced in dollars per megapixel, from $0.014 (klein 4B) to $0.10 per MP (flex); first and additional megapixels priced differently | bridge: same page: 1 credit = $0.01 (line 12) | lines 12,28-39 | sha12 9906a21a7c93 Source: https://help.bfl.ai/articles/7986977817-what-are-the-costs-associated-with-using-your-models

Datadog Bits AI (dataset row: Datadog (LLM Observability), different product replacing the dataset row)

- [D/approximate/qty yes] Datadog AI credits (pricing page) | Bits Database Optimization | ~0.3 credits | bridge: no dollar price per AI credit located on this page | lines 99-125 | sha12 d3cd03ff075c | NEAR_CAP Source: https://www.datadoghq.com/pricing/
- [D/approximate/qty yes] Datadog AI credits (pricing page) | Bits Chat message | ~0.5 credits | bridge: no dollar price per AI credit located on this page | lines 99-125 | sha12 d3cd03ff075c | NEAR_CAP Source: https://www.datadoghq.com/pricing/
- [D/approximate/qty yes] Datadog AI credits (pricing page) | Bits Agent Builder run | ~1 credit | bridge: no dollar price per AI credit located on this page | lines 99-125 | sha12 d3cd03ff075c | NEAR_CAP Source: https://www.datadoghq.com/pricing/
- [D/approximate/qty yes] Datadog AI credits (pricing page) | Bits Code fix | ~5 credits | bridge: no dollar price per AI credit located on this page | lines 99-125 | sha12 d3cd03ff075c | NEAR_CAP Source: https://www.datadoghq.com/pricing/
- [D/approximate/qty yes] Datadog AI credits (pricing page) | Bits Investigate autonomous investigation | ~4 credits | bridge: no dollar price per AI credit located on this page | lines 99-125 | sha12 d3cd03ff075c | NEAR_CAP Source: https://www.datadoghq.com/pricing/
- [D/exact/qty yes] Datadog AI credits (pricing page) | Bits Testing | free of charge through Nov 30 | bridge: no dollar price per AI credit located on this page | lines 101-103 | sha12 d3cd03ff075c | NEAR_CAP Source: https://www.datadoghq.com/pricing/

Intercom Operator (dataset row: Intercom Fin, different product replacing the dataset row)

- [P/none/qty n/a] Intercom Operator credits (Pro add-on) | Operator chat or action | page: quick question uses very little, building or testing a procedure uses a lot more; "an action costs the same number of credits for every customer"; no amounts on pages read | bridge: same page: $0.01 per credit above 2,000 included per month (lines 27, 123) | lines 152 | sha12 52a527a79a2d Source: https://www.intercom.com/help/en/articles/15841959-how-operator-credits-work-with-the-pro-add-on
- [P/none/qty n/a] Intercom Fin pricing page | Fin AI Agent | $0.99 per outcome (page line 226); not credit-metered, excluded from credit analysis | out of credit scope, excluded from the credit-action denominator | bridge: none stated | lines 226 | sha12 10d017ed52f8 | NOT COUNTED Source: https://www.intercom.com/pricing

Devin (Cognition) (dataset row: Cognition (Devin + Windsurf), different product replacing the dataset row)

- [P/none/qty n/a] Devin billing (docs) | Agent sessions (ACUs / on-demand credits) | ACU consumption described qualitatively; no ACU amounts per task on captured page | bridge: none stated | lines 12-13 | sha12 553862428596 Source: https://docs.devin.ai/admin/billing/usage

Listen Labs

- [D/exact/qty yes] Listen Labs Interview Credits (user guide) | Completed interview, self-recruited | 1 credit | bridge: no credit-to-$ rate; Platform Fee priced by commitment per SOW | lines 42 | sha12 679e4bad2903 Source: https://userguide.listenlabs.ai/
- [C/range_or_estimate/qty yes] Listen Labs Interview Credits (user guide) | Completed interview, general consumers | 5 to 8 credits | bridge: none stated | lines 42 | sha12 679e4bad2903 Source: https://userguide.listenlabs.ai/
- [C/range_or_estimate/qty yes] Listen Labs Interview Credits (user guide) | Completed interview, niche segments | 10 to 20 credits | bridge: none stated | lines 42 | sha12 679e4bad2903 Source: https://userguide.listenlabs.ai/
- [C/range_or_estimate/qty yes] Listen Labs Interview Credits (user guide) | Completed interview, targeted decision-makers | 40+ credits | bridge: none stated | lines 42 | sha12 679e4bad2903 Source: https://userguide.listenlabs.ai/

Devin (Cognition) (dataset row: Cognition (Devin + Windsurf), different product replacing the dataset row)

- [P/none/qty n/a] Devin billing (docs) | Windows vs Linux session | relative multiplier only (about 9% more than Linux); no baseline credit amount per session | bridge: none stated | lines 26 | sha12 553862428596 Source: https://docs.devin.ai/admin/billing/usage

Databricks

- [D/exact/qty no] Databricks Model Serving pricing page (GPU serving) | GPU model serving: T4 1 GPU, A10G 1/4/8 GPU, L40S 1 GPU, A100 8 GPU, RTX PRO 6000, H100 1/8 GPU | 10.48 to 800.00 DBUs per hour depending on GPU configuration | bridge: no $ per DBU on this page; a $ per DBU appears on a different product page (Lakehouse Real-Time, GCP) and is NOT applied here | lines 12-25 | sha12 40597c0d31a2 Source: https://www.databricks.com/product/pricing/model-serving
- [I/none/qty n/a] Databricks pricing page | SQL, jobs, other workloads | DBU defined; per-workload DBU rates are on SKU pages not captured | bridge: none stated | lines 82-84 | sha12 afc5400613ac Source: https://www.databricks.com/product/pricing

World Labs

- [D/exact/qty yes] World API (docs.worldlabs.ai API billing) | World generation usage event (marble-1.0, marble-1.1 fixed credits; draft) | 1,500 credits (marble-1.0, marble-1.1), 150 (draft) | bridge: same page: $1.00 per 1,250 credits (purchase rate, line 17), so 1,500 credits = $1.20 | lines 17,51 | sha12 3836ba192281 Source: https://docs.worldlabs.ai/api/pricing
- [P/none/qty n/a] Marble 2 Developer API (beta billing page) | Developer API tasks (per-task rate card) | no per-task credit amounts on the public page; it says to use the current rate card shown in the platform (line 18); beta conversion 1,000 credits per US dollar (line 14) | bridge: same page: 1,000 credits per dollar (beta); a different conversion from the World API page | lines 14,18 | sha12 cc946cd3f9bb Source: https://docs.atlas-beta.worldlabs.ai/billing
- [D/exact/qty yes] Marble Studio (beta billing page) | Frameboy session in Marble Studio (up to one minute) | 250 credits ($0.25) per session, charged once after it ends; free if under 5 seconds or the server stops responding | bridge: same page: 250 credits = $0.25 (1,000 credits per dollar, beta); the quote names a Marble Studio session, not the Developer API | lines 22 | sha12 cc946cd3f9bb Source: https://docs.atlas-beta.worldlabs.ai/billing
- [C/range_or_estimate/qty no] World API (docs.worldlabs.ai API billing) | World generation usage event (marble-1.1-plus; variable by world size) | no fixed price; observed succeeded costs through June 25, 2026: mean $1.714, median $1.280, P95 and maximum $2.480 | bridge: same page: $1.00 per 1,250 credits (purchase rate, line 17), so 1,500 credits = $1.20 | lines 96-113 | sha12 3836ba192281 Source: https://docs.worldlabs.ai/api/pricing
- [D/exact/qty yes] World API (docs.worldlabs.ai API billing) | HQ mesh export (Export API) | 3,500 credits ($2.80) per world export; one export yields both mesh variants; no charge if a mesh already exists or an export is already running | bridge: same page: $1.00 per 1,250 credits (line 17), so 3,500 credits = $2.80, as the page states | lines 17,117-120 | sha12 3836ba192281 Source: https://docs.worldlabs.ai/api/pricing

Snowflake (AI credits)

- [D/exact/qty no] Snowflake Service Consumption Table (AI Credits and Platform Credits) | AI_COMPLETE claude-4-sonnet (Table 6(a), per model) | 1.80 AI Credits input, 9.00 output, 0.180 cached input, each per one million tokens | bridge: $2.00 per AI Credit On Demand Global; $2.20 Regional; capacity tiers $1.88 to $2.00 (Table 2(b), lines 410-418 of the full capture) | lines 1339 | sha12 ef326e45e454 Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf
- [D/exact/qty user_set] Snowflake Service Consumption Table (AI Credits and Platform Credits) | AI_PARSE_DOCUMENT by mode (Table 6(g)) | Layout 3.66, OCR 0.68, AGENTIC 8.18 AI Credits per 1,000 pages | bridge: $2.00 per AI Credit On Demand Global; $2.20 Regional; capacity tiers $1.88 to $2.00 (Table 2(b), lines 410-418 of the full capture) | lines 1725-1727 | sha12 ef326e45e454 Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf
- [D/exact/qty no] Snowflake Service Consumption Table (AI Credits and Platform Credits) | Cortex Search serving (Table 6(g)) | 6.3 AI Credits per GB per month of indexed data | bridge: $2.00 per AI Credit On Demand Global; $2.20 Regional; capacity tiers $1.88 to $2.00 (Table 2(b), lines 410-418 of the full capture) | lines 1737 | sha12 ef326e45e454 Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf
- [D/exact/qty no] Snowflake Service Consumption Table (AI Credits and Platform Credits) | AI Sensitive Data Classification, openai-gpt-5-mini (Table 6(g)) | 0.15 AI Credits per one million input tokens; 1.2 per one million output tokens | bridge: $2.00 per AI Credit On Demand Global; $2.20 Regional; capacity tiers $1.88 to $2.00 (Table 2(b), lines 410-418 of the full capture) | lines 1728-1730 | sha12 ef326e45e454 Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf
- [D/exact/qty no] Snowflake Service Consumption Table (AI Credits and Platform Credits) | Batch Cortex Search (Table 6(g)) | 0.12 AI Credits per GB per hour of indexed data | bridge: $2.00 per AI Credit On Demand Global; $2.20 Regional; capacity tiers $1.88 to $2.00 (Table 2(b), lines 410-418 of the full capture) | lines 1731 | sha12 ef326e45e454 Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf
- [D/exact/qty no] Snowflake Service Consumption Table (AI Credits and Platform Credits) | Dynamic Model Routing (Table 6(g)) | 0.10 AI Credits per one million input tokens; 0.10 per one million output tokens, in addition to the underlying model price | bridge: $2.00 per AI Credit On Demand Global; $2.20 Regional; capacity tiers $1.88 to $2.00 (Table 2(b), lines 410-418 of the full capture) | lines 1738-1740 | sha12 ef326e45e454 Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf
- [D/exact/qty no] Snowflake Service Consumption Table (AI Credits and Platform Credits) | Cortex Tools (Table 6(g)) | 0 AI Credits per 1k tool calls for Snowflake-managed and third-party agents to access Snowflake endpoints; 8 AI Credits per 1k tool calls for third-party agents to access non-Snowflake endpoints | bridge: $2.00 per AI Credit On Demand Global; $2.20 Regional; capacity tiers $1.88 to $2.00 (Table 2(b), lines 410-418 of the full capture) | lines 1733-1736 | sha12 ef326e45e454 Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf

Appendix B2. RQ3 records

Each line: [level] surface | evidence | line span in the captured page | sha12 of the capture, then the source link.

Runway

- [B] Runway app | Balance on Plans and Billing page; no per-generation record described on page read | lines 15-33 | sha12 d3fd5b4d80bc Source: https://help.runwayml.com/hc/en-us/articles/15124877443219-How-do-credits-work

ElevenLabs

- [T] ElevenLabs workspace usage analytics API | API returns credit usage broken down by product type over time | lines 11 | sha12 782ab7ab3fbc Source: https://elevenlabs.io/docs/api-reference/analytics/workspace/usage

Synthesia

- [T] Synthesia billing settings | credits allocated by workspace breakdown; balance | lines 113 | sha12 ce3beaca88dd Source: https://docs.synthesia.io/docs/billing-settings

HeyGen

- [B] HeyGen credit plans | balance shown as you create; credit cost shown in-product before generating (ex ante display) | lines 170 | sha12 f7c508a15caa Source: https://help.heygen.com/en/articles/15125761-heygen-credit-based-pricing-plans-subscriptions-explained

Perplexity

- [S] Perplexity Computer, Usage page | usage breakdown by type and by session; a single task: Usage dropdown in the session shows credits used and type | lines 230-236 | sha12 d88eb4abdb35 Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity

Notion

- [T] Notion credits | Settings > Notion AI > Usage: allowance used, left, reset; owner dashboard when limits allow overage | lines 191-219 | sha12 199673d365b6 Source: https://www.notion.com/help/what-are-notion-credits

Clay

- [T] Clay Settings > Usage | dashboard of credit consumption across workspace with filters by date, owner and other fields | lines 12 | sha12 e57d8641476c Source: https://www.clay.com/faq/how-do-i-monitor-my-usage-of-data-credits-and-actions

Lovable

- [A] Lovable message history | cost of each message shown in message history (hover over the three dots); usage details split build and run credits | lines 41 | sha12 8d7903152277 Source: https://lovable.dev/pricing

Relevance AI (dataset row: not in the dataset, added, not in dataset)

- [S] Relevance AI run screen | Credits used on a run: pop-up shows each Tool cost, Agent LLM cost and base run cost | lines 73-77 | sha12 b9a2c97618b4 Source: https://relevanceai.com/docs/admin/subscriptions/plans

GitHub Copilot

- [T] GitHub Copilot usage page | usage this cycle versus budget; usage report by SKU; chart of models used | lines 15-28 | sha12 13305d3228b9 Source: https://docs.github.com/en/copilot/how-tos/manage-and-track-spending/monitor-ai-usage

Gamma

- [B] Gamma app (help center) | help FAQ: no credit history for each gamma or action; balance only | lines 142 | sha12 4c4ff9d20ddb Source: https://help.gamma.app/en/articles/7834324-how-do-credits-work-in-gamma
- [A] Gamma API | credits.deducted per agent job and credits field on generation status; page note: informational, may be omitted on lookup failure (page text: informational, not a billing record) | lines 17,97 | sha12 3f25d6aee601 Source: https://developers.gamma.app/get-started/access-and-pricing

Suno

- [B] Suno help center: How many credits do I have left | remaining credits shown in toolbar, account page and mobile profile; no per-song or usage record described on the page | lines 9-20 | sha12 01f73e568907 Source: https://help.suno.com/en/articles/2525441

Cloudflare Workers AI

- [T] Cloudflare dashboard | Neuron usage can be monitored in the Workers AI dashboard | lines 18 | sha12 a19bee9c0bee Source: https://developers.cloudflare.com/workers-ai/platform/pricing/

Snowflake (AI credits)

- [A] Snowflake ACCOUNT_USAGE view CORTEX_AI_FUNCTIONS_USAGE_HISTORY | columns: QUERY_ID, FUNCTION_NAME, MODEL_NAME, WAREHOUSE_ID, CREDITS, METRICS (tokens or pages). Attribution is per query

(job) and function and model within hourly windows, not per individual function call | lines 10-33 | sha12 a4ad8ea01def Source: https://docs.snowflake.com/en/sql-reference/account-usage/cortex_ai_functions_usage_history

Databricks

- [I] Databricks | no usage-view page captured (DBU defined only) | lines 82-84 | sha12 afc5400613ac Source: https://www.databricks.com/product/pricing

Krea

- [T] Krea Usage Statistics page | Usage Statistics page in account profile; monthly tracking; updates within 1-2 minutes; per-generation detail not stated | lines 32-45,76-78 | sha12 ed61a26151b3 Source: https://www.krea.ai/docs/user-guide/help-and-support/compute-units

Glean

- [T] Glean billing credits dashboard (help article) | dashboard: consumption over time, breakdown by feature, export | lines 10 | sha12 42b0d307237e Source: https://docs.glean.com/administration/management/billing-credits-dashboard

Genspark

- [U] Genspark /credit-usage page | balance, remaining allocation, refresh date and a usage log; per-entry credit amount not stated on page read | lines 156-162 | sha12 e040774aff42 Source: https://www.genspark.ai/helpcenter/credits-guide

World Labs

- [A] World API usage page (docs.worldlabs.ai/api/pricing) | You may view your usage event history in the usage page; each usage event may have its own cost in credits and the total cost of an API request is the sum of the usage events it maps to | lines 31 | sha12 3836ba192281 Source: https://docs.worldlabs.ai/api/pricing

Listen Labs

- [B] Listen Labs docs (email campaign) | credits deducted from account credit balance, 1 per completed response; Usage & Billing page shows balance (incentive docs); no per-study or per-interview credit record described | lines 56-66 | sha12 96bcb45a3ea1 Source: https://docs.listenlabs.ai/setup-to-launch/email-campaign

Devin (Cognition) (dataset row: Cognition (Devin + Windsurf), different product replacing the dataset row)

- [T] Devin Settings > Usage and limits | usage, quota remaining, balance; enterprise Consumption Analytics; My analytics ACU consumption per user | lines 82-87 | sha12 553862428596 Source: https://docs.devin.ai/admin/billing/usage

Salesforce Agentforce

- [I] Digital Wallet | Digital Wallet is named as the way to track Agentforce usage; the page gives no detail on what it shows | lines 83,91 | sha12 df9e9b9e067e Source: https://www.salesforce.com/agentforce/pricing/

Intercom Operator (dataset row: Intercom Fin, different product replacing the dataset row)

- [T] Intercom Usage page | credits used this period, prepaid remaining and teammate breakdown (each teammate chats and credits used) | lines 45 | sha12 52a527a79a2d Source: https://www.intercom.com/help/en/articles/15841959-how-operator-credits-work-with-the-pro-add-on

ChatGPT Business, Enterprise and Edu credit rate card (OpenAI) (dataset row: OpenAI (ChatGPT consumer), surface differs from dataset row)

- [S] ChatGPT Work and Codex chat usage details | lifetime credits per Work or Codex chat; page states it is not a complete record (sub-agents, images, background work may be missing) | lines 70-95 | sha12 48bae1c0f700 Source: https://help.openai.com/en/articles/20001478-reviewing-work-and-codex-usage-and-using-personal-analytics-in-chatgpt-desktop

Black Forest Labs

- [T] BFL dashboard | credit balance; usage breakdown by model and project in organization settings | lines 12-16 | sha12 b956bea030b9 Source: https://help.bfl.ai/articles/6872928189-where-do-i-check-my-credits

Datadog Bits AI (dataset row: Datadog (LLM Observability), different product replacing the dataset row)

- [T] Plan & Usage > AI Credits | current month consumption broken down by AI product | lines 45-47 | sha12 91e8e441ca1b Source: https://docs.datadoghq.com/account_management/billing/ai_credits/

Adobe Firefly

- [B] Adobe Firefly FAQ (first-party): Account menu counter | counter shows monthly credits allocated and credits used this cycle; no per-action record described | lines 1334-1339 | sha12 5b0afdd74bc2 | NEAR_CAP Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html

World Labs

- [I] Marble app | Marble and World API credits are separate and not transferable; Marble usage view not captured | lines 73 | sha12 9a107452ffba Source: https://docs.worldlabs.ai/marble/support/account-billing

Microsoft Copilot Credits for Copilot Studio and usage-based features (dataset row: Microsoft 365 Copilot, surface differs from dataset row)

- [S] Copilot Cowork (Microsoft Learn /cost page) | /cost shows the approximate credits for the open task so far, an aggregate for all actions in the chat session, not a line-by-line breakdown, and not an authoritative billing record | lines 8,12 | sha12 e342c9d081d0 Source: https://learn.microsoft.com/en-us/microsoft-365/copilot/usage-based-billing-copilot-credits-cost
- [T] Admin center Cost Management (Microsoft Learn) | Overview and Consumption tabs show consumption by users, groups, agents and services; Overview refreshes every 4 hours | lines 6,12,126 | sha12 a3c4d48ac0b6 Source: https://learn.microsoft.com/en-us/microsoft-365/copilot/usage-based-billing-copilot-credits-monitor-spending

World Labs

- [T] Marble 2 Developer API (beta billing page) | Usage page shows credits spent per day grouped by task or project, and by API key for Developer API usage; a day can be opened for a breakdown; Studio and Developer usage shown separately | lines 39 | sha12 cc946cd3f9bb Source: https://docs.atlas-beta.worldlabs.ai/billing

Appendix B3. Supplementary same-domain read (Nov 6, 2026)

Read separately from the frozen Nov 4 baseline, on the first-party domains of the 11 products with no stated or plan-derived dollar rate on the coded pages. No logins. Each entry gives the line numbers of the quoted text in the captured page and the first 12 hex characters of the SHA-256 of that captured text. Rates are price divided by credits as shown, for the named billing route, with arithmetic as written. Nothing here is a one-event quote, and no coded row was changed.

Adobe Firefly (R3). Firefly Standard US$9.99/mo, 2,000 generative credits per month. Route: monthly billing shown. Rate: $9.99 / 2,000 = $0.0050 per credit. Source: https://www.adobe.com/products/firefly/plans.html read 2026-11-06 PT by rendered page text; line 26: "US$9.99/mo" ; line 65: "2,000 per month", sha12 ba4172e63112. Pairs with: coded generative-credit rows (1 credit per standard generation; per-second video rows). Caveat: Standard also lists unlimited standard image features, so the price bundles features.

Clay (R3). data credits 2.5K per month at $125 per month (actions priced separately from $60 per month; Clay plan FAQ: Launch from $185 per month with 2,500

data credits). Route: monthly billing. Rate: $125 / 2,500 = $0.0500 per data credit. Source: https://www.clay.com/pricing read 2026-11-06 PT by fetch; line 136: "2.5K/ mo ($125/mo)" ; line 122: "Starts at $60/mo", sha12 cc36adafdf81. Pairs with: coded Find email and Find mobile number rows (Data Credits). Caveat: pricing page also labels 3K data credits per month; Clay plan FAQ and Clay University state 2,500 (Clay FAQ capture line 18, sha12 c6aaf1a0068c; Clay University capture line 48, sha12 2c0384ce8939), so 2,500 is used; annual figure not used.

Databricks (not counted in R3). worked example only: SKU price $0.07 per DBU (AWS, US East N. Virginia). Route: example, one cloud and region. Rate: stated example $0.07 per DBU (no division). Source: https://www.databricks.com/product/ pricing/foundation-model-serving read 2026-11-06 PT by fetch; line 16: "if your SKU price is $0.07 / DBU", sha12 b13e443211af. Pairs with: nothing counted; the coded GPU serving rows are on a different product page. Caveat: a worked example on one cloud and region for the Foundation Model Serving page, not the coded GPU serving page; the coded row states that a rate from another product page is not applied, so it is listed for completeness and not counted.

Gamma (R3). Plus 1,500 monthly credits, $108 per seat billed annually. Route: annual billing per seat. Rate: $108 / (1,500 x 12) = $0.0060 per credit. Source: https:/ /gamma.app/pricing read 2026-11-06 PT by fetch; line 32: "$108 per seat, billed annually" ; line 36: "1,500 monthly credits", sha12 43e6ec12b080. Pairs with: coded Gamma API per-card and per-image credit rows (same credit balance). Caveat: per seat; the monthly-billing price is not in the text read.

Genspark (R3). Plus 10,000 credits per month. Route: $24.99 monthly billing; $19.99 annual billing. Rate: $24.99 / 10,000 = $0.0025; $19.99 / 10,000 = $0.0020 per credit. Source: https://www.genspark.ai/business/pricing read 2026-11-06 PT by fetch; line 160: "Plus is $19.99 per month billed annually or $24.99 month to month" ; line 72: "10,000 credits per month", sha12 d215231f873d. Pairs with: coded credit guide (cost drivers only, no amounts). Caveat: plan also includes a weekly Free Quota and AI Drive storage.

Krea (R3). Basic 5,000 units per month; Krea blog: Basic $5 per month billed yearly, $9 monthly (Krea blog capture line 32, sha12 2745063c146a). Route: monthly $9; yearly $5. Rate: $9 / 5,000 = $0.0018; $5 / 5,000 = $0.0010 per unit. Source: https:/ /www.krea.ai/pricing read 2026-11-06 PT by fetch; line 88: "5,000 units / month" ; line 96: "$5/month", sha12 e9e8e46bf672. Pairs with: coded compute units per generation (docs table). Caveat: plan bundles other benefits, including unlimited relaxed generations.

Lovable (R3). Pro tiers: 100 credits $25 per month; 1,200 credits $294 per month. Route: monthly billing (annual $250 for 100 credits). Rate: $25 / 100 = $0.2500; $294 / 1,200 = $0.2450 per credit. Source: https://docs.lovable.dev/introduction/ subscription-plans read 2026-11-06 PT by fetch; line 154: "| 100 | $25 |" ; line 158: "| 1,200 | $294 |", sha12 96ab55c14f4e. Pairs with: coded Plan Mode message (1 credit). Caveat: Lovable states credit value and consumption vary by plan and feature; docs add subagent research to a Plan Mode message.

Suno (R3). Pro 2,500 credits per month; $10 monthly or $8 a month billed $96 a year. Route: monthly $10; yearly $8. Rate: $10 / 2,500 = $0.0040; $8 / 2,500

= $0.0032 per credit. Source: https://suno.com/pricing read 2026-11-06 PT by rendered page text; line 52: "Billed $96/year" ; line 57: "2,500 credits per month", sha12 8834f6c2ae6b. Pairs with: coded song generation (10 credits for two songs). Caveat: plan bundles commercial rights and downloads.

Synthesia (R3). Starter 1,250 credits per month at $19 a month (first month $9). Route: monthly billing; annual allowance not shown in the text read. Rate: $19 / 1,250 = $0.0152 per credit. Source: https://www.synthesia.io/pricing read 2026-11-06 PT by fetch; line 38: "1,250 credits / month" ; line 36: "$19 $9 first month per month", sha12 7f3ef20358fd. Pairs with: coded video generation (2 credits per second, 120 per minute). Caveat: pricing page also gives 1,000 to 2,000 credits for a one-minute agent video (a different feature); plan bundles seats and features.

Glean (no figure found). no dollar figure or credit price in the text read. Source: https://www.glean.com/pricing read 2026-11-06 PT by rendered page text; sha12 a96e68b61dcd. Caveat: not evidence that no price exists; the page offers a demo.

Listen Labs (no figure found). no dollar figure in the text read. Source: https://listenlabs.ai/pricing read 2026-11-06 PT by rendered page text; sha12 725295270b09. Caveat: user guide says Platform Fee is priced by commitment; not evidence that no price exists.

Table P. Prices and derived prices per credit unit by billing route, each product in its own unit. The units are not common and the plans bundle different features, so the rows are not a ranking or a value comparison.

Product	Credit unit	Plan and allowance	Billing route	Price	Price divided by allowance
Adobe Firefly	generative credit	Standard, 2,000 per month	monthly	$9.99	$0.0050
Clay	data credit	Price component for 2,500 per month	monthly	$125	$0.0500
Gamma	credit	Plus, 1,500 per month, per seat	annual	$108 per seat per year	$0.0060
Genspark	credit	Plus, 10,000 per month	monthly	$24.99	$0.0025
Genspark	credit	Plus, 10,000 per month	annual	$19.99 per month	$0.0020
Krea	compute unit	Basic, 5,000 per month	monthly	$9	$0.0018
Krea	compute unit	Basic, 5,000 per month	yearly	$5 per month	$0.0010
Lovable	credit	Pro, 100 per month	monthly	$25	$0.2500
Lovable	credit	Pro, 100 per month	annual	$250 per year	$0.2083
Lovable	credit	Pro, 1,200 per month	monthly	$294	$0.2450
Suno	credit	Pro, 2,500 per month	monthly	$10	$0.0040
Suno	credit	Pro, 2,500 per month	yearly	$8 per month	$0.0032
Synthesia	credit	Starter, 1,250 per month	monthly	$19	$0.0152

Appendix B4. Expiry and rollover statements located (Nov 6, 2026)

Read from the captured pages already held for the 22 products (coded pages and, for Suno and Synthesia, the Nov 6 pricing captures). Each entry quotes one located statement on expiry, reset or rollover, with the line number in the captured page and the first 12 hex characters of its SHA-256. Only whether a statement was located is coded here, not its effect on cost, and the captures are partial, so a product with no statement located may state a rule on a page not captured.

Of the 22 products, 17 have a statement located. Ten say plan credits reset, expire or do not carry over at the end of the period or term (Adobe Firefly, Genspark, Krea for individual plans, Notion, Perplexity, Runway for Standard, Pro and Unlimited, Salesforce Agentforce, Suno, Synthesia, World Labs). Seven state some non-expiry or rollover (Black Forest Labs, Clay for data credits, ElevenLabs, Gamma for purchased credits, HeyGen, Lovable for up to two months, Snowflake for contract capacity).

Adobe Firefly. Line 677: "No, generative credits do not roll over to the next month." (balance resets monthly; free-plan credits expire one month after allocation). Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html sha12 5b0afdd74bc2.

Black Forest Labs. Line 14: "Your credits don't expire" (credits also roll over to the next month). Source: https://help.bfl.ai/articles/8031234243-what-are-credits sha12 eb5542751b1c.

Clay. Line 641: "Data Credits work more like a currency and do roll over." (Actions do not roll over; data credits bank up to 2x the monthly amount on Launch and Growth). Source: https://www.clay.com/faq/what-are-clay-credits-and-how-do-they-work sha12 347c1e000f2d.

ElevenLabs. Line 214: "Unused credits roll over for up to two months". Source: https://elevenlabs.io/pricing sha12 6fe96760808d.

Gamma. Line 27: "These don't expire while your plan is active" (purchased credits; plan credits refresh and the page names a rollover cap). Source: https://help.gamma.app/en/articles/7834324-how-do-credits-work-in-gamma sha12 4c4ff9d20ddb.

Genspark. Line 171: "Credits don't carry over" (top-up credits are valid for 3 months). Source: https://www.genspark.ai/helpcenter/credits-guide sha12 e040774aff42.

HeyGen. Line 133: "Unused monthly plan credits roll over for one additional billing cycle" (on a paid plan). Source: https://help.heygen.com/en/articles/15126059-how-to-use-credits-on-heygen sha12 6dd246da1901.

Krea. Line 62: "Compute units currently do not roll over for individual plans." (business and enterprise plans roll over). Source: https://www.krea.ai/docs/user-guide/help-and-support/compute-units sha12 ed61a26151b3.

Lovable. Line 47: "monthly plan credits expire two (2) months after they're issued" (top-up credits last twelve months). Source: https://lovable.dev/pricing sha12 8d7903152277.

Notion. Line 42: "Unused credits don't carry over." Source: https://www.notion.com/help/what-are-notion-credits sha12 199673d365b6.

Perplexity. Line 249: "Unused monthly credits do not roll over." Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity sha12 d88eb4abdb35.

Runway. Line 43: "Monthly credits expire on your billing date." (Standard, Pro and Unlimited; Max may roll over up to one month). Source: https://help.runwayml.com/hc/en-us/articles/15124877443219-How-do-credits-work sha12 d3fd5b4d80bc.

Salesforce Agentforce. Line 277: "Unused Flex Credits do not roll over into subsequent subscription terms." Source: https://www.salesforce.com/agentforce/pricing/ sha12 df9e9b9e067e.

Snowflake. Line 1849: "any such Capacity Rollover will be" (contract capacity term, not a statement about AI credits alone). Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf sha12 ef326e45e454.

Suno. Line 121: "Credits included in subscriptions do not carry over from day to day or month to month." (top-up credits do not expire but need an active subscription). Source: https://suno.com/pricing sha12 8834f6c2ae6b.

Synthesia. Line 686: "Your plan's monthly allowance and credits reset each billing period." (credits bought separately are valid for 12 months). Source: https://www.synthesia.io/pricing sha12 7f3ef20358fd.

World Labs. Line 77: "do not roll over to the next month of your account billing cycle" (top-up credits differ; beta API credits do not expire). Source: https://docs.worldlabs.ai/marble/support/account-billing sha12 9a107452ffba.

No statement located in the captured pages: the five products Cloudflare Workers AI, Databricks, GitHub Copilot, Glean and Listen Labs. This does not show that no rule exists, that the term is inapplicable (for example a usage meter with no prepaid balance), or that it is undisclosed. Databricks has a trial-credit terms page that says only that trial credits are not transferable. Listen Labs has a page that advises buyers to ask about expiry and rollover, which is not a statement of its own rule.

Appendix B5. Out-of-credits and failed-run statements located (Nov 6, 2026)

Read from the captured pages already held for the 22 products. Each entry quotes one located statement, with the line number in the captured page and the first 12 hex characters of its SHA-256. Four states are kept apart: located (a statement of the rule), conditional (the page says the rule depends on a contract, a feature or a case), not applicable on the page read (the page describes billing with no prepaid allowance), and not located (no statement in the captured pages, which does not show that no rule exists or that the term is undisclosed). Only the state is coded, not the effect on cost.

What happens when the allowance runs out: 17 located, 2 conditional (ElevenLabs, Glean), 1 not applicable on the page read (Databricks), 2 not located (Black Forest Labs and Listen Labs). Snowflake's entry is a billing basis (on demand or capacity), not an overage rule.

What happens when the allowance runs out

Adobe Firefly (located). Line 85: "When you have exhausted your monthly credits, you can purchase an additional Firefly plan directly in the product". Source: https://helpx.adobe.com/africa/illustrator/using/generative-credits-access-and-use.html sha12 daf098d0560a.

Clay (located). Line 534: "One-time top-ups have a 30% premium" (Actions cannot be bought as one-time top-ups). Source: https://www.clay.com/faq/what-are-clay-credits-and-how-do-they-work sha12 347c1e000f2d.

Cloudflare Workers AI (located). Line 16: "To use more than 10,000 Neurons per day, you need to sign up for the Workers Paid plan" (further operations fail with an error when a limit is exceeded). Source: https://developers.cloudflare.com/workers-ai/platform/pricing/ sha12 a19bee9c0bee.

Databricks (not applicable). Line 14: "Databricks offers you a pay-as-you-go approach with no up-front costs" (the page describes pay-as-you-go billing with no prepaid allowance). Source: https://www.databricks.com/product/pricing sha12 afc5400613ac.

ElevenLabs (conditional). Line 52: "Additional Credits" (listed as a plan feature; the terms are not on the page). Source: https://elevenlabs.io/pricing sha12 6fe96760808d.

Gamma (located). Line 86: "If you run out of credits, you can still edit your gammas by hand." (the page also describes auto-recharge). Source: https://help.gamma.app/en/articles/7834324-how-do-credits-work-in-gamma sha12 4c4ff9d20ddb.

Genspark (located). Line 171: "Separately purchased top-up credits are valid for 3 months" (free weekly quota switches to credits when it runs out). Source: https://www.genspark.ai/helpcenter/credits-guide sha12 e040774aff42.

GitHub Copilot (located). Line 20: "additional usage is billed in GitHub AI Credits at the per-token rates". Source: https://docs.github.com/en/copilot/reference/copilot-billing/models-and-pricing sha12 16d48da537e6.

Glean (conditional). Line 80: "Your commercial contract may include additional terms (for example, minimums, overage pricing" (terms depend on the contract). Source: https://docs.glean.com/administration/management/billing-credits-dashboard sha12 42b0d307237e.

HeyGen (located). Line 39: "you can purchase additional credits or turn on auto-reload" (Business plan; Creator and Pro are told to upgrade). Source: https://help.heygen.com/en/articles/15126059-how-to-use-credits-on-heygen sha12 6dd246da1901.

Krea (located). Line 1929: "Business plans allow you to purchase additional compute packs" (free plan waits for a daily refill). Source: https://www.krea.ai/pricing sha12 8356f238dc4a.

Lovable (located). Line 283: "When your workspace has no credits left, building stops" (top-ups and auto top-up are available on Pro and Business). Source: https://docs.lovable.dev/introduction/credits-and-usage sha12 690dc51da5d1.

Notion (located). Line 58: "you can buy more mid-cycle to keep features running". Source: https://www.notion.com/help/what-are-notion-credits sha12 199673d365b6.

Perplexity (located). Line 134: "Once it runs out, tasks consume credits." (applies to included weekly usage). Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity sha12 d88eb4abdb35.

Runway (located). Line 35: "Admins can also turn on auto refill so generation continues without interruption once the monthly pool runs out." (Team plan). Source: https://help.runwayml.com/hc/en-us/articles/15124877443219-How-do-credits-work sha12 d3fd5b4d80bc.

Salesforce Agentforce (located). Line 289: "There is no overage penalty." (exceeding the entitlement is billed at the contracted rate monthly in arrears). Source: https://www.salesforce.com/agentforce/pricing/ sha12 df9e9b9e067e.

Snowflake (located). Line 295: "On Demand basis, where usage is invoiced in arrears every month" (or a Capacity basis purchased up front; a billing basis, not a stated overage rule). Source: https://www.snowflake.com/legal-files/CreditConsumptionTable.pdf sha12 ef326e45e454.

Suno (located). Line 16: "you can purchase additional credits by visiting suno.com/account" (a 50-credit daily allotment applies after monthly credits run out). Source: https://help.suno.com/en/articles/2643713 sha12 110ba29e4cbc.

Synthesia (located). Line 45: "you are able to top-up mid-cycle". Source: https://help.synthesia.io/en/articles/8364334-how-do-i-check-my-credit-balance-in-synthesia sha12 525cd6cdddac.

World Labs (located). Line 81: "All paid plans allow you to purchase additional top-up credits at any time when you run out.". Source: https://docs.worldlabs.ai/marble/support/account-billing sha12 9a107452ffba.

Failed runs: 3 located (Adobe Firefly, Clay, Gamma), 3 conditional (Genspark, Perplexity, World Labs), 16 not located (Black Forest Labs, Cloudflare Workers AI, Databricks, ElevenLabs, GitHub Copilot, Glean, HeyGen, Krea, Listen Labs, Lovable, Notion, Runway, Salesforce Agentforce, Snowflake, Suno, Synthesia). No statement was inferred from silence.

Failed runs

Adobe Firefly (located). Line 928: "On an unsuccessful generation, credits are typically refunded to your credit balance within minutes." (refund can take up to 24 hours for longer jobs such as video). Source: https://helpx.adobe.com/creative-cloud/apps/generative-ai/generative-credits-faq.html sha12 5b0afdd74bc2.

Clay (located). Line 85: "Failed lookups do not consume credits.". Source: https://agent-docs.clay.com/faq/enrichment-platform-usage-log-finance-spend-by-user sha12 eeaf1dc2b3ae.

Gamma (located). Line 16: "You're not charged credits when an AI action fails". Source: https://help.gamma.app/en/articles/7834324-how-do-credits-work-in-gamma sha12 4c4ff9d20ddb.

Genspark (conditional). Line 202: "your credits are refunded automatically" (only when a prompt is blocked by the safety filter before anything is generated; a completed generation that is later flagged is still charged, and each failed agent retry costs credits). Source: https://www.genspark.ai/helpcenter/credits-guide sha12 e040774aff42.

Perplexity (conditional). Line 333: "credits are not automatically refunded" (when a task finishes but the result is not what the user wanted; a task that cannot finish because of an error on the provider side is referred to support review). Source: https://www.perplexity.ai/help-center/en/articles/13838041-how-credits-work-on-perplexity sha12 d88eb4abdb35.

World Labs (conditional). Line 22: "A session is free if it plays for less than 5 seconds" (one Studio feature only). Source: https://docs.atlas-beta.worldlabs.ai/billing sha12 cc946cd3f9bb.