LINDY

LINDY · LINDY.AI

IN THIS TEARDOWN

01 The model
02 The meter
03 What the page commits you to
04 Three changes that would move revenue
Evidence · Edges

PREPARED FOR
public read

READ DATE
2026-10-17

PRICING PAGE
https://www.lindy.ai/pricing

INDEX CATEGORY
agent-platform

COMPARISON SET
6 vendors

SCOPE
public data only

EVERY FIGURE SOURCED · ~ MARKS A COMPUTED FIGURE · NOT A RANKING · NOT ADVICE ON WHAT TO CHARGE

THE READ

Lindy sells one paid plan, Team, as a per-user credit pack on a 15-step ladder from a $29.99 floor to $9,999.50 a month, with top-ups at a flat $10 per 1,000 credits. On 2026-10-17 the page served two different floors to the same reader: 7 of 10 loads showed Team at $29.99 for 3,000 credits, 3 of 10 showed $49.99 for 5,000, same $0.01 per credit on both, same ladder above 5,000. Lindy is testing how much a new customer must commit, not what a credit costs. The structure costs Lindy in two places: the ladder's volume discount, 43% from floor to top, sits inside a closed dropdown that no search snippet, text read or first glance can see, and two steps on that ladder (20,000 and 45,000 credits) cost more per credit than the step below them. The largest lever is to put the ladder's rate on the page and fix the two steps, so the discount the buyer already gets starts pulling upgrades instead of top-ups. This read replaces the 2026-10-13 memo, which saw one floor and no ladder.

PART 01

01 THE MODEL

Hybrid: a per-user subscription that buys a monthly credit allotment, chosen from a dropdown ladder, plus top-up credits at a published rate. OBSERVED

Fit verdict: the credit ladder fits an agent platform; the seat does not, and the entry test says Lindy knows the floor is the problem. ANALYSIS The footnote bills every human ("Every active user in your workspace is billed, from $29.99 a month for 3,000 credits each") while the work pools ("Every seat adds its credits to a shared pool"). A seat meter on an agent product prices against the product's own success: the better Lindy works, the fewer humans need to be active in it. The ladder is where value is counted, and the ladder is hidden.

What changed since the Index row (verified 2026-09-18: Plus $29.99, Pro $99.99, Max $199.99 per user per month): the three tiers are gone and one card, Team, carries a credit ladder. The old $29.99 survives as one of two floors under test; the old $99.99 and $199.99 survive as the 15,000- and 35,000-credit steps. OBSERVED The entry price did not rise 67%, as the 2026-10-13 read concluded from a single load; one of two floors is 67% higher than the other, with 67% more credits. ANALYSIS

The test itself: ten reads between 04:19 and 04:25 UTC on 2026-10-17, one text fetch and nine browser loads with a fresh session each. Seven served the $29.99 / 3,000 card, three served $49.99 / 5,000. The site's navigation also varies between loads ("Hire Lindy for" against "Solutions") independently of the price shown, so at least two experiments run on the page at once. The meta description reads "Starts at $29.99 per user per month" on every load. OBSERVED

Category context (Index v1.1, 6 agent-platform rows, verified 2026-09-08 to 2026-09-18): ANALYSIS

- 6 of 6 publish a price. 6 of 6 gate a tier behind sales. Lindy is at parity on both.
- 5 of 6 are usage-credits. Lindy is the only hybrid, the only vendor in the set that bills per human.
- Four peers publish a USD entry price: Make.com $12, Relevance AI $29, Zapier Agents ↯ $33.33 (annual-only, $400 ÷ 12), Gumloop $37. n8n Cloud publishes €20 only. Lindy's $29.99 floor sits inside that range; its $49.99 floor sits above every peer that publishes a USD entry.
- 3 of 6 publish an annual discount (Relevance AI 33%, n8n 17%, Make.com 15%). Lindy publishes none and offers no annual billing on the page.
- 3 of 6 have a perpetual free tier. Lindy's "Free" is "$50 in credits" that "last 7 days": a trial with a free label.
- Three vendors call the unit a credit. Price per credit at entry: Make.com ↯ $0.0012, Gumloop $0.005, Lindy ↯ $0.01. The credits are different units of work, so this compares prices and says nothing about value per credit.

PART 02

02 THE METER

BILLABLE UNIT credit. IN-PLAN UNIT PRICE ↯ $0.01 at the floor ($29.99 ÷ 3,000 or $49.99 ÷ 5,000), falling to ↯ $0.0057 at 35,000 credits and above. TOP-UP / OVERAGE $0.01 ($10 ÷ 1,000). SPREAD ↯ 1.0x at the floor, ↯ 1.75x from 35,000 credits up.
COUNTABLE BEFORE PURCHASE no. RISES WITH VALUE partly. VARIANCE CARRIED BY buyer on usage, vendor on model cost.

1. **Countable before purchase: no.** The page publishes three estimate bands, "2-250 Credits", "250-1,000 Credits", "1,000-2,500 Credits", and no rate per action or per model; "Credits measure the work Lindy does" is the whole definition. OBSERVED The bands span 1,250x. In dollars at the floor rate they are ↯ $0.02 to $2.50, $2.50 to $10, and $10 to $25 per task. The $29.99 floor buys ↯ 1.2 of the largest task class a month; the $49.99 floor buys 2. ANALYSIS
2. **Rises with value: partly.** A big task costs more than a small one, so cost tracks size. It does not track whether the task succeeded; a build that took three tries costs three builds. ANALYSIS
3. **Spread: 1.0x at the floor, 1.75x at scale, and the ladder is not monotonic.** At the floor, top-up credits cost what in-plan credits cost, and monthly credits expire while top-ups do not, so there is no reason to hold the allotment. From 15,000 credits the in-plan rate falls below the top-up rate and the subscription starts earning its place. Two steps break the curve: 20,000 credits at $149.99 is ↯ $7.50 per 1,000 against ↯ $6.67 at the 15,000 step below it, and 45,000 at $299.99 is ↯ $6.67 against ↯ $5.71 at 35,000. ANALYSIS For comparison, Relevance AI's top-up is ↯ 6.9x its in-plan rate (Index row, verified 2026-09-18). OBSERVED
4. **Variance.** Usage variance sits on the buyer, with a stop-loss: at zero credits Lindy pauses credit-using actions rather than billing. Model-cost variance sits on Lindy unless credit cost varies by model, which the page does not say. ANALYSIS

For the sales motion: the floor is a clean procurement decision and the ladder is a real volume discount that no buyer sees until they open a dropdown on the card. Every expansion conversation today is "buy top-ups at list" or "Book a demo", while a 43% discount sits unadvertised between them. ANALYSIS

PART 03

03 WHAT THE PAGE COMMITS YOU TO

> Two floors on one URL, "$29.99/mo" for "3,000 credits / mo" and "$49.99/mo" for "5,000 credits / mo", served at random → Whichever wins, a share of this month's signups saw the other one. The losing arm's buyers hold a screenshot of a price you offered.

> "Starts at $29.99 per user per month" in the meta description → Every search snippet promises the lower floor. Keeping the $49.99 floor means every buyer who arrives from search reads a 67% increase before the FAQ.

"Top up anytime, $10 per 1,000 credits" → $0.01 is the public price of a credit. A top-up above it reads as an increase; a floor-step cut reads as a cut to list.

A 15-step ladder behind "Buy the credits you need." → You own a published volume discount, ↯ 43% from floor to top, and two steps that charge more per credit than the step below them. A buyer who finds the kink distrusts the rest of the ladder.

"Every active user in your workspace is billed, from $29.99 a month for 3,000 credits each" beside a dropdown labelled "Monthly credits for your team" → You owe an answer to whether a step is per user or per workspace. The page gives both.

Credit-using actions pause at zero, "Nothing breaks and there is no surprise bill" → A hard stop. A workflow that runs out on day 20 goes dark until an admin acts, and revenue per seat is capped at the step chosen.

PART 04

04 THREE CHANGES THAT WOULD MOVE REVENUE

1 Decide the floor, put the ladder's rate on the page, and fix the two steps that cost more per credit than the one below

RECOMMENDATION

MECHANISM: Expansion first, conversion second. The ladder is a 43% volume discount no buyer sees before clicking a dropdown; printing "$10 per 1,000 at the floor, $5.71 at 35,000 and up" makes it the upgrade path that top-ups and "Book a demo" absorb today. Fixing the 20,000 and 45,000 steps removes the two a buyer should never pick. Closing the floor test takes the second price out of every search snippet.

EVIDENCE: Two floors served to the same reader, 7 of 10 and 3 of 10 OBSERVED. Ladders of 15 and 14 steps, identical from 5,000 credits up except a 7,000 step OBSERVED. Unit rates ↯ $10.00, $6.67, $7.50, $5.71, $6.67, $5.71 per 1,000 at 8,000 through 70,000 credits ANALYSIS. Three of six peers (Relevance AI, n8n Cloud, Make.com) publish a priced tier above entry on the page itself ANALYSIS.

SIZE: Direction up on upgrades from the floor and on average step chosen. The size needs two of Lindy's numbers: conversion and 60-day retention by floor arm, and the distribution of credits consumed per workspace against the step chosen.

RISK: Printing the ladder rate shows floor buyers they pay 75% more per credit than large accounts, and the fixed steps cut revenue from anyone sitting on 20,000 or 45,000 today. Signal: floor-tier churn rising in the first 60 days while upgrades do not. Keep the floor-to-top spread where it is, under 2x.

2 Replace the hard pause with an admin-set auto-top-up ceiling, defaulted on at checkout

RECOMMENDATION

MECHANISM: Expansion, with churn as the second effect. Every seat that hits zero mid-cycle is a revenue event that becomes a pause. An admin who sets a ceiling at checkout ("top up to this amount per month, then pause") keeps the no-surprise-bill promise and converts the pause into a top-up.

EVIDENCE: Credit-using actions pause at zero and "Nothing breaks and there is no surprise bill" OBSERVED. The top-up rate is published, "Top up anytime, $10 per 1,000 credits", so the ceiling is a number the buyer can read before they buy OBSERVED. At the $29.99 floor the pause arrives after ϟ 1.2 large tasks ANALYSIS.

SIZE: Direction up on top-up revenue, down on mid-cycle churn of scheduled workflows. The size needs the share of seats that reach zero before cycle end and how many days early.

RISK: A ceiling defaulted too high, or hidden in settings, breaks the promise the page makes twice. Signal: refund requests and support tickets mentioning unexpected charges in the first billing cycle after launch.

3 Publish the exchange rate: credits per task class by model

RECOMMENDATION

MECHANISM: Conversion. The buyer who can forecast the bill picks a ladder step; the one who cannot picks the floor, books a demo, or leaves. A rate card by model closes the gap between "2-2,500 credits" and a number procurement can budget, and it is what makes change 1's ladder choosable.

EVIDENCE: Three bands and no rate per model OBSERVED. In the category, Make.com defines its credit exactly and prices it at ϟ $0.0012; Gumloop defines its credit at $0.005 (Index rows verified 2026-09-08 and 2026-09-18) OBSERVED. Lindy's $0.01 floor credit is the most expensive credit in the category and the least defined ANALYSIS.

SIZE: Direction up on trial-to-paid conversion and on average step chosen. The size needs Lindy's trial conversion rate and the stated reasons for non-conversion.

RISK: Publishing the rate exposes one of two facts: expensive models cost more credits, so buyers route to the cheapest model; or they do not, and the page now documents that Lindy absorbs the model-cost difference. Signal: average credits per task falling after publication while task count holds.

EVIDENCE

#	CLAIM	QUOTE (VERBATIM)	SOURCE URL	READ DATE	LABEL
1	Free plan	"$0"; "$50 in credits for free"; "Free credits last 7 days"	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
2	Team card, arm A	"$29.99" "/mo"; "3,000 credits / mo"	https://www.lindy.ai/pricing (headless Chromium, 6 of 9 loads; text fetch 1 of 1)	2026-10-17	OBSERVED
3	Team card, arm B	"$49.99" "/mo"; "5,000 credits / mo"	https://www.lindy.ai/pricing (headless Chromium, 3 of 9 loads)	2026-10-17	OBSERVED
4	Footnote, arm A	"Every active user in your workspace is billed, from $29.99 a month for 3,000 credits each."	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
5	Footnote, arm B	"Every active user in your workspace is billed, from $49.99 a month for 5,000 credits each."	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
6	Ladder, arm A (15 steps)	"3,000 credits" "$29.99"; "4,000 credits" "$39.99"; "5,000 credits" "$49.99"; "6,000 credits" "$59.99"; "8,000 credits" "$79.99"; "15,000 credits" "$99.99"; "20,000 credits" "$149.99"; "35,000 credits" "$199.99"; "45,000 credits" "$299.99"; "SMALL TEAMS" "70,000 credits" "$399.98"; "105,000 credits" "$599.97"; "175,000 credits" "$999.95"; "GROWING TEAMS" "350,000 credits" "$1,999.90"; "700,000 credits" "$3,999.80"; "1,750,000 credits" "$9,999.50"	https://www.lindy.ai/pricing, dropdown opened	2026-10-17	OBSERVED
7	Ladder, arm B (14 steps)	"5,000 credits" "$49.99"; "6,000 credits" "$59.99"; "7,000 credits" "$69.99"; "8,000 credits" "$79.99"; then identical to arm A from "15,000 credits" "$99.99" to "1,750,000 credits" "$9,999.50"	https://www.lindy.ai/pricing, dropdown opened	2026-10-17	OBSERVED
8	Below the ladder	"Need more? Top up anytime at $10 per 1,000 credits, or talk to sales."	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
9	Top-up rate on the card	"Top up anytime, $10 per 1,000 credits"	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
10	Dropdown label	"Buy the credits you need."; "Monthly credits for your team"	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
11	Credits pool	"Every seat adds its credits to a shared pool, so anyone can draw on it when they need more."	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
12	Enterprise	"Custom"; "Custom credits and billing"; "Book a demo"	https://www.lindy.ai/pricing	2026-10-17	OBSERVED

#	CLAIM	QUOTE (VERBATIM)	SOURCE URL	READ DATE	LABEL
13	Credit definition and bands	"Credits measure the work Lindy does. You only use them when Lindy is working."; "2-250 Credits"; "250-1,000 Credits"; "1,000-2,500 Credits"	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
14	Expiry	monthly credits do not roll over; top-up credits stay available until used (FAQ answers, text fetch)	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
15	Pause at zero, no overage	credit-using actions pause until the next cycle or an admin tops up; "Nothing breaks and there is no surprise bill." (FAQ answer, text fetch)	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
16	Meta description, both arms	"Starts at $29.99 per user per month." (meta description and og:description, every load)	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
17	Navigation varies by load	"Hire Lindy for" on 3 loads, "Solutions" on 6, each label seen with both price arms	https://www.lindy.ai/pricing	2026-10-17	OBSERVED
18	Read tally	10 reads 04:19 to 04:25 UTC: 1 text fetch (arm A) + 9 headless-Chromium loads, fresh session each (6 arm A, 3 arm B); derived from the logged loads	this memo's read log	2026-10-17	ANALYSIS
19	Prior reads of the same page	2026-10-13 text read: $49.99 / 5,000. 2026-10-14, two sessions: $49.99 / 5,000. 2026-10-15, two sessions five minutes apart: $29.99 / 3,000 then $49.99 / 5,000	Ventures/PricingIndex/teardowns/index-corrections.md	2026-10-17	REPORTED
20	Index Lindy row	"$29.99 /mo per user"; entry 29.99 (Plus), mid 99.99 (Pro), top 199.99 (Max); unit_price_usd 0.01 derived; verified_date 2026-09-18	https://ignatiustheyoungerai.com/pricing-index/agent-pricing-index.json	2026-10-17	OBSERVED
21	Category size and shares	6 rows category "agent-platform"; publishes_a_price true 6 of 6; enterprise_contact_sales true 6 of 6; pricing_model usage-credits 5, hybrid 1	same dataset, v1.1, last_verified 2026-09-18	2026-10-17	ANALYSIS
22	Peer entry prices	Make.com 12 (2026-09-08); Relevance AI 29 (2026-09-18); Zapier Agents 33.33 derived from "$400 billed annually" (2026-09-08); Gumloop 37 (2026-09-18); n8n Cloud €20, no USD (2026-09-18)	same dataset	2026-10-17	OBSERVED
23	Peer unit prices	Make.com 0.0012 derived; Gumloop 0.005 stated; Relevance AI 0.0116 derived; Zapier Agents 0.0222 derived	same dataset	2026-10-17	OBSERVED
24	Relevance AI top-up spread	notable_terms: top-ups $80 per 1,000 Actions, about 6.9x the in-plan rate	same dataset	2026-10-17	OBSERVED
25	Annual discounts and free tiers in category	annual_discount_pct: Relevance AI 33, n8n Cloud 17, Make.com 15; null for Lindy, Gumloop, Zapier Agents. Perpetual free tier: Relevance AI, Zapier Agents, Make.com	same dataset	2026-10-17	OBSERVED
26	Peers with a priced tier above entry	Relevance AI (29 to 349), n8n Cloud (€20 to €667), Make.com (12 to 38); Gumloop and Zapier Agents top = entry	same dataset	2026-10-17	ANALYSIS

#	CLAIM	QUOTE (VERBATIM)	SOURCE URL	READ DATE	LABEL
27	In-plan unit price at the floors	↯ $29.99 ÷ 3,000 = $0.009997; $49.99 ÷ 5,000 = $0.009998; both ↯ $10.00 per 1,000	rows 2, 3	2026-10-17	ANALYSIS
28	In-plan rate along the ladder	↯ per 1,000 credits: $10.00 at 3,000 through 8,000; $6.67 at 15,000; $7.50 at 20,000; $5.71 at 35,000; $6.67 at 45,000; $5.71 at 70,000 through 1,750,000	row 6	2026-10-17	ANALYSIS
29	Spread	↯ $0.01 ÷ $0.009997 = 1.0x at the floor; $0.01 ÷ $0.005714 = 1.75x from 35,000 credits up (and from 70,000 without exception)	rows 9, 27, 28	2026-10-17	ANALYSIS
30	Ladder discount, floor to top	↯ 1 - ($9,999.50 ÷ 1,750,000) ÷ ($29.99 ÷ 3,000) = 42.8%, stated as 43%	row 6	2026-10-17	ANALYSIS
31	The two kinks	↯ 20,000 credits at $149.99 ($7.50 per 1,000) above 15,000 at $99.99 ($6.67); 45,000 at $299.99 ($6.67) above 35,000 at $199.99 ($5.71). $50.00 more buys 15,000 more credits at the 35,000 step; $99.99 more buys 25,000 more at the 70,000 step	row 6	2026-10-17	ANALYSIS
32	Floor difference	↯ $49.99 ÷ $29.99 - 1 = 66.7%; 5,000 ÷ 3,000 - 1 = 66.7%	rows 2, 3	2026-10-17	ANALYSIS
33	Bands in dollars and large tasks per floor	↯ 2–250 cr = $0.02–$2.50; 250–1,000 cr = $2.50–$10; 1,000–2,500 cr = $10–$25; 3,000 ÷ 2,500 = 1.2 and 5,000 ÷ 2,500 = 2	rows 9, 13	2026-10-17	ANALYSIS
34	Free credits in units	↯ $50 ÷ $0.01 = 5,000 credits, one month at the $49.99 floor or 1.67 months at the $29.99 floor	rows 1, 9	2026-10-17	ANALYSIS

EDGES

- **This read replaces the 2026-10-13 memo.** That memo read the page once, by text, saw the $49.99 / 5,000 card and no dropdown, and concluded the entry price had risen 67% from the Index row. On 2026-10-17 the page served $29.99 / 3,000 to 7 of 10 reads and $49.99 / 5,000 to 3. The earlier conclusion is withdrawn; the 67% is the gap between two floors under test, not a change in list price. The 2026-10-13 PDF should not be cited for Lindy's entry price.
- **Which arm is "the price" cannot be determined from the page.** Both floors were served within the same minute to fresh sessions. This memo reports both and picks neither. A buyer, a journalist or an Index row quoting one number is quoting half the page.
- **The ladder was read from an opened dropdown in a headless browser,** not from the page's text layer, which does not contain it. A text-only reader (search engines, agents fetching Markdown) sees one price per load and no ladder.

- **Per user or per workspace is unresolved.** The footnote bills every active user "for N credits each"; the dropdown is labelled "Monthly credits for your team"; the FAQ says credits pool. Whether a 35,000-credit step is per seat or per workspace changes the invoice by the seat count, and the page gives both readings.
- **Not read today:** FAQ answer text beyond what the text fetch returned (collapsed in the browser), whether credit cost varies by model, the Enterprise price, what "active user" means, whether the 7,000 step exists on arm A under some condition.
- **Peer rows were not re-read today.** Zapier Agents and Make.com were verified 2026-09-08, the rest 2026-09-18. Corrections queued since (Make.com annual discount, Perplexity, Agentforce) are not applied to the dataset this memo cites.
- **No crowd median exists for this category.** The aggregates endpoint was not re-read today; the 2026-10-13 read held 2 submissions, both in "other". Nothing from it is cited.
- **The anchor is missing.** The ceiling on Lindy's price is the cost of the human who would do the work the bands describe, against ↯ $2.50 to $25 per task. Neither the page nor the Index gives that human cost. It is the single number that would most change Part 4.
- **What Lindy's numbers would sharpen:** conversion and 60-day retention by floor arm (decides change 1's floor); the distribution of credits consumed per workspace against the step chosen (sizes change 1 and sets which kink matters); gross margin per credit by model (sizes change 3 and the safe ladder slope); the share of seats that reach zero before cycle end (sizes change 2); trial-to-paid conversion and reasons for non-conversion (sizes change 3); top-up share of revenue (sets the cannibalization risk in change 1).

IGNATIUSTHEYOUNGERAI · Pricing teardown of Lindy · Read date 2026-10-17 · Comparison set drawn from the Agent Pricing Index, ignatiustheyoungerai.com/pricing-index · Labels: OBSERVED is what a source states; ANALYSIS is a read across sources; MODELED is computed from stated assumptions; RECOMMENDATION is a proposed change. Corrections: hello@ignatiustheyoungerai.com