

# How Kashi prices, and why 可視

Pricing is the single-biggest unforced-error surface in B2B SaaS. Undercharging burns capital. Overcharging kills conversion. Picking the wrong pricing metric makes both happen at the same time. This document explains the four frameworks B2B SaaS companies actually use to set prices, applies them to Kashi, and lays out a 90-day plan to validate the numbers before locking them.

Drafted 2026-04-21 · Companion to [/business.html](#) §4 + [/jp-launch-runbook.html](#)

## 4 frameworks

value-based, cost-plus, competitive, hybrid — most SaaS uses hybrid

## 3 tiers

Good / Better / Best + Enterprise — the industry standard shape

## ~80%

share of SaaS revenue the middle tier typically captures

## 10 conversations

customer WTP interviews before locking any provisional price

## 00 Why pricing is strategy, not math

There are three pricing failure modes that kill B2B SaaS companies before they reach scale. In descending order of how often they actually happen:

1. **Wrong pricing metric.** The company charges per-seat but the value scales with something else (meetings, data volume, governance depth). Every customer feels the price is misaligned with what they get. Churn is high for reasons the team can't name.
2. **Undercharging.** The founders are self-conscious about asking for money, they anchor to their own salary-adjacent intuition, they set the price at 1/5 of what the buyer would gladly pay. Every deal feels fine until the runway runs out.
3. **Freemium-without-conversion-path.** A generous free tier draws users; none of them ever upgrade because the free tier covers their real need. Free tier becomes an expensive charity program.

Each of these mistakes is avoidable with a deliberate process. "Pick a number that feels right" is the process almost every first-time founder uses, and it is the process that produces these three failures. This doc is the alternative.

**The single sentence that defines good B2B pricing.** The price should be a fraction of the economic value the customer captures from using the product, measured along a metric that scales with that value.

## 01 The four canonical B2B SaaS pricing frameworks

Every B2B SaaS pricing decision uses one of these four, or a hybrid of them. Understanding which one applies to Kashi is the prerequisite to picking actual numbers.

### FRAMEWORK 1 **Value-based pricing** USE FOR KASHI

Price as a fraction of the economic value the customer captures from the product. The gold standard for B2B software because marginal costs are near-zero and the only defensible price ceiling is customer willingness-to-pay, not cost.

**Typical capture rate:** B2B SaaS companies usually charge 10-25% of the value they deliver. Anything higher risks buyer reluctance. Anything lower leaves money on the table.

**For Kashi:** value = averted cases of quiet resignation and mental-health leave, plus productivity recovery from earlier intervention. The rebuilt-ROI memo puts a 500-person company at ¥131M/year of productivity drag from team-dynamics issues. A 10-20% reduction via earlier intervention is ¥13-26M/year recovered. Charging 10-20% of that value gives ¥1.3-5.2M/year per 500-person company. Current Professional-tier pricing at ¥800/user/mo lands at ¥4.8M/year for a 500-person company — right in the middle of the value-capture band. Math confirms the provisional price.

### FRAMEWORK 2 **Cost-plus pricing** AVOID FOR SAAS

Price = COGS + target gross margin. How restaurants, retailers, and most physical-goods companies price. Almost always wrong for software.

**Why it fails for SaaS:** the marginal cost of serving one more user is near-zero. A cost-plus calculation produces a floor, not a ceiling, and following it strictly leads to race-to-the-bottom pricing. Cost-plus still matters as a *check* (your gross margin must be healthy — for SaaS, 70-85% is the target band), but not as the primary method.

**For Kashi:** treat cost-plus as a gross-margin sanity check. Starter tier gross margin is N/A (free). Pro tier gross margin targets 85% (compute + email support). Enterprise tier gross margin targets 65% because the CSM + consultation hours are meaningful cost. These are healthy bands. If any tier drops under 60% gross margin, the price is too low or the tier is overscoped.

### FRAMEWORK 3 **Competitive anchoring** SANITY CHECK ONLY

Price as a function of the closest competitor's price  $\pm 20-50\%$ . Useful when the market has established reference prices the buyer already knows. Dangerous as a primary method because it implicitly concedes that your product is a substitute for the competitor, which undercuts differentiation.

**For Kashi:** the closest category (engagement survey: Wevox, Geppo, Culture Amp) prices at ¥300-1,200/user/mo. The deeper category (harassment-scanning: Archaic, FRONTEO) prices higher but doesn't publish. Kashi's current Pro tier (¥800/user/mo) sits at the top of the engagement-survey band and below the harassment-scanning band. Defensible as "governance infrastructure, priced above surveys because it does more than surveys, priced below content-scanners because we refuse to read content." Use this framing in sales, but do not let it set the number. The number comes from value-based math.

### FRAMEWORK 4 **Hybrid (industry standard)** KASHI'S ACTUAL MODEL

Most B2B SaaS companies combine all three. Different parts of the product get priced differently.

- **Bottom of funnel (freemium / Starter):** priced near zero to drive adoption. Cost-plus is relevant because you don't want the free tier to bankrupt you.
- **Middle (Professional / mid-market):** competitive-anchored, because mid-market buyers do compare shop. Priced as "better than the obvious alternative" with a clear

differentiation story.

- **Top (Enterprise):** value-based, custom-negotiated, often not published. The value of Enterprise features (governance-compliance, SLAs, SOC 2, residency) is asymmetric across buyers, so pricing has to flex.

**For Kashi:** this is already the model. Starter free (bottom), Professional per-user (middle, anchored against surveys + content-scanners), Enterprise base + per-user (top, value-based, negotiated). The decision tree that matters is not whether to hybrid, but what goes in each tier.

## 02 Choosing a pricing metric (what you charge FOR)

The pricing metric matters more than the price. A good metric makes every customer feel the cost matches their use. A bad metric produces arguments about every invoice.

METRIC	BEST WHEN...	BREAKS WHEN...	FIT FOR KASHI
<b>Per-seat (user)</b>	Value scales ~linearly with # of people using the product. Buyer already counts seats.	Light users subsidize heavy users. Buyers hit "seat ceilings" and find ways to share logins.	<b>Pro tier</b> — value scales with # of employees whose meetings flow through the detectors.
<b>Per-meeting analyzed</b>	Usage is highly variable per customer. Cost tracks closely with usage.	Budget-predictability erodes. Finance teams dislike consumption billing for governance tools.	<b>Poor fit</b> — governance budget is planned annually; variable billing feels wrong for the category.
<b>Per-team (bundle)</b>	Natural usage unit is the team, not the individual.	Large companies with many small teams pay too much; flat-team pricing caps upside.	<b>Weak fit</b> — team boundaries are fluid; seat-based pricing captures the same value more flexibly.
<b>Flat-rate per</b>	You want simplicity. SMBs love it.	Leaves enterprise money on the table; a	<b>Starter only</b> — fine for the ≤20-employee free tier. Wrong

METRIC	BEST WHEN...	BREAKS WHEN...	FIT FOR KASHI
<b>company</b>		50-person co pays the same as a 5,000-person co.	for revenue.
<b>Base + per-seat (hybrid)</b>	Enterprise value has fixed + variable components. Consultation + certification drive fixed; per-user drives variable.	Small buyers find the base intimidating; you need a clear story for why the base exists.	<b>Enterprise tier</b> — current Enterprise uses ¥10M base + ¥800/user/mo. The base pays for hands-on consultation, CSM, 就業規則 review, SOC 2 access — not for the software.

**Kashi's metric recommendation.** Confirmed: per-seat for Professional, base+per-seat for Enterprise, flat (free) for Starter. The current pattern already aligns with industry best-practice for governance-category SaaS.

### 03 Tier strategy — Good / Better / Best + Enterprise

The classic B2B SaaS tier pattern exists for behavioral reasons, not mathematical ones.

- **Entry tier** — freemium or very cheap. Purpose: get users in the door, let them experience the product, create a conversion funnel. Typically 10-30% of total users, 0-5% of revenue.
- **Middle tier ("the workhorse")** — the landing zone. Priced to feel "reasonable for a serious team." Typically 40-60% of users, 60-80% of revenue. This is where pricing discipline matters most.
- **Top tier** — the anchor. Exists as much to make the middle tier look reasonable as to serve real customers. Usually 5-15% of users, 20-35% of revenue.
- **Custom Enterprise** — above the top tier, not published, negotiated. Revenue share varies wildly by category; for governance-compliance SaaS it's often 40-60% of total.

The trick is feature-gating — deciding which capabilities belong to which tier. Kashi's current feature gates (proposed):

ENTRY

WORKHORSE (LANDING)

## STARTER

### Free

≤20 employees

- 3 core structural detectors
- Self-visibility only
- No executive view
- 90-day hard-delete retention
- Community support

## PROFESSIONAL

### ¥800/user/mo

≤500 employees

- All 7 detectors (incl. keigo asymmetry)
- Manager Mirror + Executive Brief
- Feedforward commitments
- 24-month retention
- JP data residency (Tokyo primary)
- Downloadable labor-consultation packet
- Email support (24h response)

## ANCHOR

### ENTERPRISE

### ¥10M/yr base

+ ¥800/user/mo · 500+ users

- Everything in Professional
- Evidence vault (E2E encrypted)
- SSO / SAML / SCIM
- Dedicated CSM
- Hands-on labor consultation (就業規則 review + rep-process sit-in)
- SOC 2 Type II + ISO 27001 report access
- 24/7 support + SLA
- NAQ-R validation study access

## NEGOTIATED

### ENTERPRISE+

### Custom

multi-region · multi-tenant

- Everything in Enterprise
- Multi-region residency (EU + US)
- Dedicated single-tenant instance
- Custom detector extensions
- Named engineering support contact
- Co-authored case-law review
- Priced via ROI-share model, not per-user

**Tier-design trap to avoid.** Do not put the keigo-asymmetry detector behind the Enterprise tier. It is Kashi's category-defining differentiator and must be visible from the Professional tier for the sales pitch to land. Keep it in Pro; use governance depth (consultation, residency, certifications, SLAs) as the Enterprise gate.

## 04 Competitor benchmarks

These are the reference prices in the Japanese workplace-analytics market. Source column lists where the number comes from — published pricing pages, customer reports, or public press.

Treat all non-published numbers as estimates within  $\pm 30\%$ .

PRODUCT	CATEGORY	PRICE	SOURCE / NOTE
<b>Geppo</b> (Recruit × CyberAgent)	Engagement survey (monthly 3-question)	¥298/user/mo from	Published on geppo.jp pricing page (2025)
<b>Wevox</b> (Atrae)	Engagement survey + ONA analytics	¥400-700/user/mo	Published list price; volume-discount bands per Atrae IR disclosures
<b>Culture Amp</b>	Engagement + performance surveys	¥800-1,200/user/mo	Industry analyst estimates (G2, TrustRadius) for JP SMB tier
<b>MS Viva Insights + Glint</b>	M365 telemetry + engagement survey	~¥5,000/user/mo (bundled in M365 E5) or ¥1,000-1,500/user/mo standalone	MS published pricing; Viva Suite SKU list
<b>Workday Peakon</b>	Engagement survey, sentiment AI	~¥1,000-1,600/user/mo	Typical Workday enterprise quote band (not published)
<b>Archaic</b> ハラスメントチェック AI	Content-scanning harassment triage	~¥1,000-2,000/user/mo (est.)	No published price; estimated from similar JP enterprise-SaaS bands
<b>FRONTEO KIBIT Eye</b>	Content-scanning email/chat triage	Enterprise only · custom	Historically deployed at MUFG / Aeon; pricing not public. Category activity has gone quiet post-2022.
<b>15Five + Kona</b>	Manager-effectiveness coaching	~\$14-20/user/mo (~¥2,100-3,000)	15Five published pricing; Kona is an add-on
<b>Kashi (Pro)</b>	Governance infrastructure (structural + keigo)	¥800/user/mo (provisional)	Current — this doc reviews whether the number holds

PRODUCT	CATEGORY	PRICE	SOURCE / NOTE
<b>Kashi (Enterprise)</b>	Governance + consultation + certifications	¥10M/yr base + ¥800/user/mo (provisional)	Current — blended ACV ~¥19.6M @ 1,000 users

## Positioning map

On a price-vs-capability plot, Kashi sits above engagement surveys (we do more than ask people how they feel) and below content-scanners (we refuse to read message bodies). The Pro tier at ¥800/user/mo is at the top of the engagement-survey band, priced by the value we deliver beyond surveys, not by matching what surveys charge.

**The sales one-liner that follows from this table:** "We charge roughly the same as Culture Amp or Wevox at the top of their band, because we do something those products do not do. We charge less than Archaic or Viva because we refuse to do something those products do — specifically, read content."

## 05 Applying the frameworks to Kashi — the recommendation

Putting frameworks §01-§04 together, the current pricing numbers are well-supported. The table below traces each tier back to its framework rationale.

TIER	PRICE	VALUE-BASED RATIONALE	COMPETITIVE RATIONALE	CONFIDENCE
<b>Starter</b>	Free · ≤20 employees	Value capture < ¥5M/yr at this scale; cost of serving one more free user is near-zero.	Free tier standard in the category (Geppo offers free trial, Wevox offers partial free).	High — keep as-is.
<b>Professional</b>	¥800/user/mo · ≤500	A 500-person co has ~¥131M/yr of team-dynamics drag. 10-20% recovery = ¥13-	Matches top of engagement-survey band. Below content-scanner floor.	Medium-high — validate with 10 WTP conversations before locking.

TIER	PRICE	VALUE-BASED RATIONALE	COMPETITIVE RATIONALE	CONFIDENCE
		26M. Charging ¥4.8M/yr = 18-37% value capture. Just below the "resist" band ( $\geq 40\%$ value capture risks buyer reluctance).	Above 15Five/Kona.	
<b>Enterprise</b>	¥10M base + ¥800/user/mo · 500+	A 1,000-person co has ~¥262M/yr of drag. ¥19.6M/yr ACV = 7.5% value capture — low end, which is correct because Enterprise also pays for consultation + certifications that have separate costs.	Not directly comparable (no competitor bundles consultation at this tier). The base is defensible on service cost alone.	Medium — Enterprise base is the most negotiable. Be prepared to flex 20-50% on the base in exchange for reference-customer status, case-study rights, or multi-year commit.
<b>Enterprise+</b>	Custom · ROI-share	At multi-region / multi-region scale, value capture is company-specific. ROI-share pricing (e.g., 15% of measured productivity recovery) aligns incentive better than flat seat pricing.	No direct competitor — this is the "governance infrastructure" claim.	Exploratory — do not publish. Negotiate when the first candidate emerges.

## Price-insensitivity bands

What happens if Pro goes to ¥1,200/user/mo (50% higher)?

- Value-capture jumps to 28-55%. The 55% end crosses the "resist" threshold — some buyers will balk on principle even if the math works.

- Competitive position shifts from "top of survey band" to "peer of Culture Amp" — fine, but weakens the "we're the reasonable option vs content-scanners" line.
- Works better if paired with an explicit value statement on the pricing page ("this replaces ¥X/year of Y"). Not recommended without that reframing.

What happens if Pro goes to ¥500/user/mo (38% lower)?

- Value-capture drops to 11–23%. Healthier for buyers, worse for unit economics (CAC payback extends from 3.3 months to 5.3 months at the same CAC).
- Moves Kashi into the middle of the engagement-survey band, implicitly conceding that we're an engagement tool. Weakens positioning.
- Recommended only if WTP research shows strong price resistance at ¥800 from the core target segment.

**Final recommendation for this pass.** Keep Professional at ¥800/user/mo and Enterprise at ¥10M base + ¥800/user/mo for the first 10 paying customers. Collect WTP data during those conversations. Revisit pricing after customer #10 or at Month 12, whichever comes first. Do not change prices in the meantime — pricing instability damages credibility more than a few percent of margin optimization recovers.

## 06 Bootstrap-phase pricing — the first 0→5 customers

The frameworks above assume you have willing-to-pay data and market leverage. In Phase 1 (the pre-revenue bootstrap) you have neither. This is the section the rest of the industry treats as an afterthought.

### The "no data, no leverage" problem

When you have zero paying customers, you face a chicken-and-egg: you can't price intelligently without WTP data, and you can't collect WTP data without running pricing conversations. Most founders guess a number. Better approach:

1. **Run 10 structured WTP conversations before putting any price anywhere.** Use Van Westendorp (see §07). Output: a band of acceptable prices for your core target segment.
2. **Pick a provisional price in the upper half of that band.** Founder psychology biases low; consciously correct upward.

3. **Put that price on the site as "Professional" and leave Starter free.** Do not publish Enterprise — say "Enterprise available; contact us."
4. **Offer the first 3-5 pilot customers a founding-customer discount** (see below).
5. **Charge customer #6 full list price.** If you don't have 5 paying customers by Month 6, the problem is product or positioning, not pricing.

### **Founding-customer pricing — cheaper than pilot concessions**

The JP-launch runbook §05 includes a ¥2-4M "pilot concession" line. For pre-revenue bootstrap, that line is **zero** because you have no revenue to concede. But you still want to compensate the first customer for the risk they take on an unproven vendor. The right instrument is *founding-customer pricing*:

- **Founding-customer price** — first 3-5 customers get a ~50% discount off provisional Pro pricing (so ¥400/user/mo instead of ¥800), locked in for 2 years. After that they pay full rate.
- **Reference-rights** — in exchange, they agree to be a named reference, provide a quote, and attend 1-2 reference calls per quarter.
- **Why this beats "free pilot"**: any money invoiced is validation; ¥0 invoiced is "they're not sure enough to charge anything." A small real price is more credible than a free trial.

**Specific bootstrap numbers (recommended).** First paid customer: ¥400/user/mo for a 2-year founding-customer window, capped at ¥500k/year for ≤100-person companies. That's roughly ¥40k/month of revenue from a 100-person pilot — meaningful psychologically (first real MRR) and easy on their budget (under the approval authority of most JP SMB CEOs).

### **Why not free?**

Free pilots look attractive — zero friction, zero commercial negotiation. But there are three reasons to charge something:

- **Intent signal.** Paid customers show up to meetings. Free-pilot users ghost when priorities shift.
- **Budget legitimacy.** If the pilot lives in "miscellaneous tools," the decision to continue is soft. If it lives in an actual budget line, the decision is hard.
- **Internal narrative.** The pilot champion at the customer company has to explain why they're spending money on this. That explanation builds organizational alignment that a free pilot never generates.

Starter remains free as a top-of-funnel surface. That's different from giving the first Professional customer a free pilot.

## 07 Willingness-to-pay research — Van Westendorp

Before locking any price, run the Van Westendorp Price Sensitivity Meter. It's a 4-question survey that produces a price-acceptance band without asking customers to directly state what they'd pay (which is unreliable).

### The four questions

1. **Too cheap:** "At what monthly per-user price would you think the product is so cheap that its quality must be suspect?"
2. **Cheap (a bargain):** "At what price would you consider the product a bargain — a great deal for the money?"
3. **Expensive (but still considering):** "At what price would the product start to feel expensive, but you'd still consider buying it?"
4. **Too expensive:** "At what price is it too expensive to consider?"

Ask each in JPY per user per month. Collect from 10+ respondents. Plot the four curves; the intersections give you four prices:

- **Point of Marginal Cheapness (PMC)** — where "too cheap" and "expensive" cross. Below this, buyers doubt quality.
- **Point of Marginal Expensiveness (PME)** — where "too expensive" and "cheap" cross. Above this, buyers refuse on principle.
- **Indifference Price Point (IPP)** — where "cheap" and "expensive" cross. The median buyer's sweet spot.
- **Optimal Price Point (OPP)** — where "too cheap" and "too expensive" cross. The price at which the fewest buyers reject on price.

Defensible price range = [PMC, PME]. Pick within that range based on positioning goals.

### Practical plan for Kashi

Van Westendorp formally requires a structured survey with 50+ respondents. Kashi doesn't have that yet. Run it informally:

1. Build the 4 questions into your Phase-1 pilot-conversation script. Ask them toward the end of the conversation once the prospect understands what the product does.
2. Collect responses from the first 10 pilot conversations. Do not share the current ¥800 price beforehand — it anchors their answers.
3. At conversation #10, plot the curves in a spreadsheet. If the current ¥800/user/mo is inside [PMC, PME], hold. If it's below PMC or above PME, revise before customer #6.

Don't spend money on formal research (SurveyMonkey Audience, Qualtrics) until after the ¥30-50M angel round. The informal version with real prospects is more predictive at this stage than a formal panel of random respondents.

## 08 Six B2B SaaS pricing mistakes, ranked

The six traps that kill early-stage B2B SaaS pricing, ordered by how often they actually happen and how much damage they do.

### 1. Pricing from COGS

Most common failure

Founder calculates "it costs us ¥50/user, add 3x margin, charge ¥150/user." In SaaS, marginal cost is near-zero. This calculation anchors the price at 1/10 of what the buyer would gladly pay.

*Mitigation:* treat COGS as a gross-margin sanity check only (keep margin >70%). Set price from value-based reasoning. See §01 Framework 1.

### 2. Wrong pricing metric

Charging per-seat when value scales with meetings; charging per-meeting when buyers plan governance budgets annually. Produces invoice arguments and churn for reasons the team can't name.

*Mitigation:* explicit metric-fit analysis in §02. Ask prospects how they budget for comparable tools; match the metric to that mental model.

### 3. Publishing Enterprise prices

Prospect sees the number, skips the conversation, self-disqualifies. Removes the negotiation ceiling that justifies high-touch Enterprise ACV.

*Mitigation:* publish Starter and Professional. Enterprise = "contact us." Exception: publish an example ACV band (¥15-25M) with a note that actual pricing depends on scope.

<b>4. Changing prices too often</b>	Pricing instability in the first 2 years signals "we don't know what we're doing." Kills trust with prospects who saw the old price, and kills renewal conversations with customers on the old price.	<i>Mitigation:</i> commit to a 12-month pricing freeze from first publication. If you must change, grandfather existing customers and notify 90 days in advance.
<b>5. Free tier without conversion path</b>	Starter covers the actual need of 90% of free users, so none ever upgrade. Free tier becomes an expensive charity.	<i>Mitigation:</i> Kashi's Starter $\leq 20$ employees is defensible because the value at that scale is genuinely small (no Mirror, no Exec Brief). Monitor: if Starter users who hit the 20-employee cap convert at $<15\%$ within 6 months, retune the gating.
<b>6. Permanent founding-customer discount</b>	"I'll give you 50% off forever" — the single worst concession in early-stage pricing. Locks in a below-cost price as revenue scales.	<i>Mitigation:</i> founding-customer discount is time-boxed (2 years) and explicit in the contract. After the window, auto-steps up to list price. Never "forever."

## 09 90-day pricing validation plan

Concrete actions, week by week, to move from "guessed number" to "defended number."

- Week 1-2 — Draft pricing page + Van Westendorp script**  
Draft a `/pricing` page with Starter (free) + Professional (¥800/user/mo) + Enterprise (contact-us). Do not publish yet. Write the 4 Van Westendorp questions into the standard pilot-conversation script.

---

- Week 3-4 — Run 5 pilot conversations with WTP questions**  
Target: CEO or CHRO at 50-300-person JP companies (see jp-launch-runbook.html §02). Run the 4 questions toward the end of each conversation. Record the numbers.

---

- Week 5-6 — Run 5 more conversations**  
Get to 10 data points. Plot the four Van Westendorp curves. Identify the [PMC, PME] band.

---

- Week 7 — Decision point**  
If ¥800/user/mo is inside [PMC, PME], publish the pricing page. If it's outside, revise (and update business.html §4 + deck.html slide 10). Lock for 12 months.

### **Week 8-12 — First pilot commercial conversations**

3 pilot candidates see the published pricing. Founding-customer discount (¥400/user/mo for 2 years) is the lever. Measure: how many balk at ¥400, how many at ¥800, how many at "Enterprise = contact us."

### **Month 3-6 — First paid customer signs**

Realized price for the first signed customer is the most important data point. Compare to the provisional Pro price. If realized > provisional, the provisional is too low. If realized < provisional after negotiation down, the provisional may be at the right level for buyers to feel like they negotiated well.

### **Month 12 — Formal pricing review**

With 5-10 paying customers, re-run the analysis. This is where the ¥30-50M angel round has typically closed, so formal Van Westendorp research (50+ respondents, paid panel) becomes affordable. Revise prices if data warrants; grandfather existing customers.

## 10 **What this doc does not cover**

- **Multi-currency pricing.** USD, EUR, SGD, GBP pricing for non-JP markets. Different purchasing-power parity, different competitive bands. Separate pass after JP pricing is validated.
- **Discount policies.** Volume discounts, annual prepay discounts, multi-year commit discounts, non-profit discounts. Requires a formal policy doc; premature before first paid customer.
- **Subsidy-aware pricing.** IT導入補助金, 働き方改革推進支援助成金, and similar JP subsidies can cover 50-75% of the customer-side cost. Full treatment lives in </jip-launch-runbook.html> §01 non-dilutive grant section.
- **Contract terms.** MSA, SLA, TOS, data processing agreement. Legal workstream. Pricing is the easier half of the commercial conversation.
- **Packaging experiments.** Bundling keigo detector separately, metered evidence-vault storage, usage caps. These are Phase 2+ optimizations after the base three-tier model is validated.

**See also:** </business.html> (§4 pricing table in the context of the full business plan) · </jip-launch-runbook.html> (90-day JP launch incl. pilot concession logic for post-revenue phase) · </research/> (rebuilt-ROI-math memo + category-entry memo both inform §01 and §05)

Drafted 2026-04-21 · Status: draft for internal review · Next pass after 10 Van Westendorp conversations or first paid customer, whichever comes first