

Purchasing power. *For every personal agent.*

The consumer-side transaction network that turns intent and flexibility into better offers.

• Public draft • 07 Sep 2026

Stackbroker is building the consumer-side transaction network for personal agents. A consumer declares a need and acceptable flexibility once. Their agent uses Stackbroker to find executable offers from opted-in providers, accept within an explicit mandate, and obtain a provider-confirmed booking and an accountable receipt. Compatible demand can later unlock offers that one consumer could not obtain alone.

Proposed financing: \$2 million pre-seed for 18 months. Public discussion draft; final financing amount and terms remain to be determined. Founder: Doug Butdorf. [Talk to Doug about investing](#), or contact [Lando](#), AI Chief of Staff.

1. The investment thesis

Businesses have sophisticated pricing and distribution systems; consumers still spend time searching fragmented availability and negotiating individually. Agents can make consumer intent persistent, portable and actionable. Stackbroker does not need to win the assistant interface: it aims to become a useful transaction capability inside many assistants.

Three advantages could compound: broader agent distribution brings credible demand; compatible demand attracts better provider offers; reliable outcomes increase repeat use. The network effect is local and conditional on service, geography, time and consumer requirements. A million unrelated requests do not create liquidity for a particular appointment. The initial product must help one consumer before it can claim collective purchasing power.

The venture case is cross-category transaction volume with low variable cost. It is not a subscription business, an enterprise governance product, a restaurant reservation reseller, or a sale of consumer data. A single-city booking concierge is a validation environment, not the venture-scale endpoint.

2. Why now, and what the timing does not prove

Scheduling APIs, MCP/REST distribution and machine-payment protocols provide practical building blocks. Square documents buyer- and seller-permissioned booking capabilities. Cal.com documents API v2, including OAuth and API-key paths. Stripe documents machine payments using MPP/cards and stablecoins, with x402 support. These are source-verified platform capabilities, not Stackbroker partnerships or enabled production integrations. [S1-S3]

Harry Stebbings wrote, "Network Effects at the Agent Level Will Separate AI Winners," in his September 7 post summarizing a conversation with Town's founder. That is industry commentary, chiefly about multi-user assistant workflows—not evidence that consumers will pay Stackbroker. The post also argues that learning from users cannot be accelerated like code generation. We take the timing signal seriously and the customer-learning constraint equally seriously. [S4]

We will use deterministic matching and workflow orchestration for routine transactions. The personal agent does expensive reasoning where necessary; Stackbroker need not finance a frontier-model inference on every quote.

3. Product: from intent to an executable offer

An illustrative request: "Haircut next week after 4 p.m.; Tuesday through Thursday; total below \$45." The consumer's agent supplies permitted substitutions, expiry, approved total, fee authority and confirmation rules. A provider can offer a \$40 appointment in a quieter period; the proposed \$0.50 booking fee makes the approved total \$40.50. This is an example, not live inventory or a savings claim.

The workflow is: declare mandate → discover or wait for an eligible offer → reserve only when the provider supports a real hold → confirm in the provider system → reconcile the fee → return a receipt. A persistent request replaces repeated agent searches when providers supply events. Where event coverage is incomplete, refreshes must be authorized, bounded and account-rate-limited.

Version one uses posted offers and standing orders. Version two adds independently priced provider offers for deduplicated, opt-in demand cohorts. No highest-bidder scarcity auction is planned. Suppliers cannot pay for ranking, see individual maximum budgets or coordinate confidential price floors through the network.

The consumer experience remains inside their assistant. A lightweight approval or receipt page covers hosts unable to present consent or detailed terms. No destination marketplace, recurring subscription, hidden commitment or forced crypto wallet is required.

4. Beachhead and category sequence

Start with one local pocket and one supported scheduling connector. Our lower-friction technical-track hypothesis is salons, barbers or auto detailing: standardized appointment types, understandable terms and merchant-controlled capacity. Provider selection must be based on actual access and customer demand; no metro or provider is presently committed.

Dental cancellation filling is an early commercial validation track, particularly existing patients seeking an earlier practice-approved routine appointment. Its value can be access rather than price. Practice eligibility, clinical suitability and insurance decisions stay with the practice. Cal.com and Square are not universal dental-system connectors. Identifiable appointment data may be sensitive health information; clinical data does not enter the shared demand book. An opaque reference is not automatically anonymous. Validate the actual practice integration and data roles before live handling.

After repeatable appointment value: extend to other scheduled services, then bounded recurring household-service negotiation. Restaurants come last because inventory access, scarcity, operations and applicable reservation rules complicate entry. Consumer-service payments initially remain with the provider; Stackbroker collects only its separately authorized success fee.

5. Market opportunity: size the transactions we can earn

We do not use total consumer-service spending as Stackbroker revenue TAM. A \$100 service is not \$100 of our revenue. The initial fee hypothesis is \$0.50 per provider-confirmed booking, and addressability depends on permissioned supply, agent adoption, transaction frequency and willingness to pay.

Illustrative annual gross fee opportunity at two bookings per monthly transacting consumer per month:

- $100,000 \text{ consumers} \times 2 \times \$0.50 \times 12 = \textbf{\$1.2 million.}$
- $1 \text{ million consumers} \times 2 \times \$0.50 \times 12 = \textbf{\$12 million.}$
- $10 \text{ million consumers} \times 2 \times \$0.50 \times 12 = \textbf{\$120 million.}$

These are scale sensitivities, not current users, a TAM estimate, ARR, net revenue or a forecast. At one transaction a month the values halve. Appointment frequency alone may not support two monthly transactions; broadening categories must earn that frequency.

The initial 10,000 monthly transacting-consumer financing target implies 20,000 bookings/month. At 40 incremental bookings per active provider/month, it requires roughly 500 active providers. Both the booking productivity and provider count are hypotheses. This exposes the operational scale needed; three initial providers are a learning cohort, not enough supply for the financing target.

Before Series A claims, build a sourced market model by category and geography using active provider counts, eligible capacity, permitted integration coverage and observed frequency. Count buyers, not agents or keys. Census nonemployer data can help scope independent service businesses but must not be conflated with bookable providers. [S6]

6. Go-to-market: density before breadth

Seed both sides together: three opted-in providers, their permissioned waitlists, and two independent personal-agent builders serving real consumers. Agent builders receive a working sandbox, clear tool descriptors, examples, fee semantics, narrow scopes and receipts—not an enterprise subscription pitch. Provider onboarding emphasizes incremental profitable capacity and no provider fee.

Pilot recruitment will test two consumer messages: “Get an earlier appointment without checking again” and “Trade flexibility for a better all-in offer.” Preserve a baseline for existing booking behavior so discounts and rescheduling can be distinguished from incremental value.

Acquisition sequence: opted-in waitlist need → real agent request → provider-confirmed booking → completed service → repeat paid request → adjacent provider invitations. Registries, llms.txt and tool stores enable discovery but do not themselves create demand. No claimed distribution partnership exists today.

Expansion is earned one repeatable local pocket at a time. An agent integration may distribute software widely, but supply activation, support and regional constraints remain real work. Measure activation cost per provider and per transacting consumer, not installation counts. Broad paid consumer advertising is unlikely to work on a fifty-cent fee until repeat economics justify it.

7. Business model and unit economics

The initial experiment is a **\$0.50 consumer-authorized fee on a provider-confirmed booking**, included in the approved total. Free bounded discovery, standing requests, status, cancellation-tool calls and failed matches. Providers pay nothing; no subscriptions, paid placement or data sales. Platform sponsorship can ease trial but unsponsored, repeat fees must prove willingness to pay.

Reverse our fee when provider or system failure defeats the booking. Define voluntary cancellation treatment before the consumer accepts. Provider deposits and cancellation penalties are separate, visible terms. Confirmed bookings, paid fees and completed services remain distinct ledger events. Revenue recognition and refunds must follow the actual service obligation, not an API response.

Stripe's ordinary US domestic online-card pricing is 2.9% + \$0.30. A \$0.50 fee leaves **\$0.1855** before infrastructure, refunds, support or acquisition. This is an ordinary-card stress case, not verified MPP/SPT account pricing. Stripe documents a \$0.50 card-SPT minimum and 0.01 USDC stablecoin minimum. Stablecoin access requires approval and excludes New York businesses under the cited Stripe availability policy. Entity location and account enablement remain unverified. [S3, S5]

Planning unit case, all assumptions rather than measured costs: \$0.50 fee less \$0.04 payment cost, \$0.03 orchestration/free-request allocation and \$0.08 refund/support allowance leaves \$0.35 contribution before acquisition. Allocate a further \$0.05 per booking to distribution/activation and contribution is **\$0.30, or 60%**, before fixed operating expense. Under ordinary-card processing with the same other allocations, contribution falls to **\$0.0255**, or 5.1%. A \$0.25 payment would be uneconomic on the ordinary card tariff even before other costs.

At the target \$0.35 pre-acquisition contribution and two monthly bookings, a consumer produces \$0.70 monthly contribution before churn. A \$2 acquisition cost would take about 2.9 months to recover; this is not measured LTV. Quotas and abuse controls must allocate unsuccessful searches to successful outcomes. One support-heavy failure can erase many micro-fees.

Do not build a custody business to solve processor fees. Investigate eligible machine-payment rails and transparently authorized aggregation where available, without transforming the offer into a subscription. Choose the rail from verified all-in costs, failure handling and consumer availability. If none supports sustainable pricing, change the fee or transaction category before expanding.

8. Competition and defensibility

Scheduling systems such as Square and Cal.com own provider workflows and can add agent access themselves. Marketplaces such as Booksy, Fresha and OpenTable own discovery relationships and supply networks. Personal-agent platforms and payment providers can also move into orchestration. These are competitive categories and examples, not a claim that any lacks AI or negotiation capabilities.

Stackbroker's proposed position is portable consumer mandates across providers, independently competing offers and reliable success-based execution. Start as an interoperable demand source rather than replacing the provider's system of record.

Defensibility must be earned through repeat distribution, permissioned supply access, reliable reconciliation, permission-scoped outcome history and local demand liquidity. MCP wrappers and generic matching code are replicable. Consumers and providers will multi-home. A genuine moat exists only if participating improves their outcomes enough to sustain repeat use; proprietary ownership of their intent is not the customer promise.

9. Product foundation and delivery architecture

The live Stackbroker foundation includes MCP/REST discovery and routing, consent-based pilot enrollment, provider intake, policy controls and permission-scoped receipts. These capabilities support distribution and onboarding; booking execution and consumer payments remain in development.

The next release connects private consumer mandates to provider-authorized availability and reliable booking fulfillment. Consumer identity, provider permissions, payment authority and transaction state remain separated. Only information needed to fulfill the request is shared, with revocable permissions and protected provider connections.

Flow: authorized inventory event → validated offer → eligible principal-scoped order → expiring match → capability-aware confirmation → provider booking result → fee/reversal reconciliation → keyed receipt. Use durable state transitions, idempotency keys and an outbox/saga pattern for external effects. Payment and booking cannot be made one distributed atomic transaction merely by using a database lock. Unknown provider results require reconciliation before retry, not a second booking.

Enforce total-cost limits, revocation, provider opt-in and account-wide rate budgets at execution. Deduplicate one principal's equivalent need across agents and cancel alternatives after fulfillment. An internal lock is not an external reservation. Booking confirmation, fee collection and service completion are recorded as separate outcomes.

Cal.com API access and OAuth eligibility require provider-specific verification; its managed-user Platform is closed to new signups. Square buyer-level permissions have limited calendar and webhook visibility; seller-level writes require Plus/Premium. Prefer supported OAuth for scaled onboarding. Choose one connector from an explicit capability matrix. [S1, S2]

10. Evidence and validation gates

As of September 7, 2026, pilot enrollment is live, with **no verified customer revenue or live consumer bookings**. Agent tools and provider intake are available on Stackbroker. The next milestone is a permissioned, provider-confirmed appointment followed by a completed service and repeat paid use. Live enrollment demonstrates product availability, not customer adoption.

Proposed early evidence bar: three real providers, two independent external agents serving real people, twenty confirmed bookings, ten completed services and some unsponsored, affirmatively authorized fees. These are minimum learning gates, not statistically validated product-market fit.

Instrument request-to-offer, offer-to-booking, booking-to-completion, fee-paid and refund rates; 30/60/90-day repeat paid use; time-to-appointment; all-in consumer savings where comparable; provider incremental contribution and discount cannibalization; support minutes; acquisition cost; and variable contribution after failed requests. Suppress sparse cohorts in demand reports.

Require a genuinely cohort-dependent offer before claiming bargaining power. If people value access but do not change times, withdraw the demand-flattening claim rather than force the narrative. If consumers refuse the fee, or all eligible rails lose money, stop the relevant mechanism. If no provider supplies executable capacity in a bounded category test, change access/category before continuing broad acquisition.

11. Milestone and operating plan

First 30 days, target sequence contingent on access: select one market/connector; secure three opt-ins; complete permission and fee-economics checks; establish baseline data; verify durable orders and failure handling. Days 31–60: deliver the initial twenty confirmed/ten completed-service learning cohort with real authorized fees. Days 61–90: observe repeat paid use, measure contribution and test pooling only if compatible density exists.

Proposed 18-month financing objectives: two production connectors; three local pockets with a reproducible activation playbook; 10,000 monthly transacting consumers as an ambitious exit target; and positive variable contribution plus observed repeat paid cohorts. The target is not a committed forecast. Expansion depends on readiness and evidence, not elapsed calendar time.

Founder Doug Butdorf leads company direction and fundraising. AI-assisted development supports rapid iteration while people remain accountable for product quality and customer outcomes. The proposed team adds a founding engineer and an integrations/full-stack engineer, with fractional design, security, legal and provider-operations support. Lando, Doug's AI Chief of Staff, supports coordination and investor inquiries.

12. Capital plan and financial sensitivities

Proposed raise: **\$2 million**, funding **18 months**. Product and engineering: \$1.10m (55%). Distribution and provider activation: \$400k (20%). Reliability, security and operations: \$300k (15%). Contingency: \$200k (10%). Total: \$2m. Deployment excluding reserve averages \$100k/month; including reserve, the budget envelope is about \$111k/month. Payroll, founder compensation, contractors and infrastructure must fit these categories without double counting.

Cost boundaries: the \$100k fixed monthly budget covers salaried/retained engineering, integration development, base infrastructure, distribution staff and fixed operations. The separate \$0.20 per-booking allowance covers incremental processor charges (\$0.04), marginal infrastructure and unsuccessful requests (\$0.03), fee-refund reserve (\$0.01), incremental support (\$0.07), and attributable distribution/activation (\$0.05). These are incremental costs, not a second allocation of the fixed payroll. The \$2 acquisition-cost example is an independent sensitivity, not an additional model expense or a cost inferred from the \$0.05 allowance.

Cash conventions: assume the entire proposed raise arrives before month one and collected booking fees settle in the same month; no financing fees, taxes, receivables, interest or separate provider-service cash flows are modeled. Modeled fee refunds reduce gross fees to net collected fees; all other variable cash expenses then reduce contribution. Real settlement delays, disputes and tax obligations can change this timing.

The companion CSV models an illustrative ramp from no users to 10,000 monthly transacting consumers in month 18; 2 bookings/consumer/month; a \$0.50 fee; \$0.20 variable cost/booking; and \$100k fixed monthly operating budget. Base and card-stress cases reach twenty bookings in month two, matching the day-60 learning target; downside can miss this gate. Monthly transacting consumers are a stock, not acquisition or retention cohorts; churn is not separately estimated before cohort evidence exists. Implied provider requirements use the same forty-booking productivity assumption and at least three pilot providers. It tracks gross fees, modeled refunds, net collected fees, contribution and cash—not ARR or GMV. Downside and upside variants change demand, not magically improve margins. No revenue is needed to claim an 18-month base runway because \$1.8m operating deployment plus \$200k reserve is already funded by the proposed raise; transaction losses could shorten it.

At the base exit target, gross fees are only \$10k/month and target contribution is \$6k/month against the \$100k monthly operating plan. The financing buys evidence and repeatability, not break-even. At \$0.30 contribution per booking, covering \$100k monthly fixed costs requires approximately 333,334 bookings/month, or 166,667 monthly transacting consumers at two bookings each. Under the ordinary-card stress case, that becomes about 3.92m bookings/month, or 1.96m monthly transacting consumers. This is why a low-cost payment path and repeat distribution are prerequisites to the venture case.

Investor return potential therefore depends on earning distribution at millions of consumers, increasing useful transaction frequency across categories and maintaining positive contribution. No valuation multiple, exponential-growth promise or financing commitment is assumed.

13. Key risks and decisions

Cold start: use existing opted-in waitlists and bounded local supply. Platform dependence: permissioned connectors, portable contracts and more than one distribution route. Thin fees: validate actual processing/support economics before scale. Weak retention: measure paid cohorts, not free orders. Agent abuse: principal deduplication and quotas. Booking errors: durable idempotent orchestration, reconciliation and an accountable human escalation path. Sensitive appointments: segregated data flows and category-specific review. Conflicted pricing: no provider-paid ranking or consumer-budget disclosure. Incumbent response: prove useful cross-provider outcomes, not a feature novelty claim.

The remaining commercial choices are evidence-driven market selection, rail eligibility, first connector and customer-approved fee terms. Expansion will follow demonstrated customer value, repeat use and viable unit economics.

Sources and assumptions

Sources checked September 7, 2026. Official documentation establishes capabilities, not access grants or partnerships. Product status is management-reported as of the date of this public draft; detailed diligence materials are available on request.

- **S1:** [Square Bookings overview](#) — permission scopes, booking operations and plan limitations.
- **S2:** [Cal.com API v2](#) — auth paths and Platform signup closure.
- **S3:** [Stripe machine payments](#) — MPP/SPT/x402 methods, minima and availability; account eligibility unknown.
- **S4:** [Harry Stebbings post](#) — retrieved post text via public syndication, not investment endorsement or independently verified funding claims.
- **S5:** [Stripe pricing](#) — ordinary US domestic-card benchmark; not an assertion of our machine-payment tariff.
- **S6:** [US Census Nonemployer Statistics](#) — potential market-research input, not used here to manufacture a TAM.

All pricing experiments, market scenarios, hiring allocations, milestone dates and financial-model values above are management-planning assumptions. They are deliberately separate from verified historical facts.

Investor contact

Talk to Doug about investing using the Outbranch contact form. For coordination, contact [Lando, AI Chief of Staff](#) — lando@outbranch.net.

[Talk to Doug about investing ↗](#)

Public draft
07 Sep 2026

[Lando · AI Chief of Staff · lando@outbranch.net](#)

[Financial model CSV](#) ↓

[Investor deck](#)

[One-page brief](#)

[Download plan PDF](#) ↓