Own the Data. Own the Category.

A Research-Driven Demand Generation Program for Zenskar

zenskar

Prepared by John Mecke, Managing Director, DevelopmentCorporate LLC

July 2026 · Confidential

Contents

Executive Summary

Zenskar has built one of the most sophisticated content operations of any Series A company in the billing and monetization category. Buyer's guides, comparison hubs, and category explainers engineered for both search and LLM synthesis have earned Zenskar challenger-slot mentions in a category whose incumbent citation slots are otherwise locked by Zuora, Chargebee, Stripe, and Recurly. That is a genuine achievement — and it has hit a ceiling.

The ceiling is epistemic authority. Every incumbent that reliably occupies a citation slot in this category publishes original primary research: Recurly's State of Subscriptions is in its tenth annual edition, Chargebee fields its State of Recurring Revenue & Monetization, and Metronome's State of Usage-Based Pricing defined the usage-billing conversation until Stripe acquired the company. Zenskar publishes none. In the AI-mediated research environment where Zenskar's own buyers — finance leaders at venture-backed companies between roughly $5M and $100M ARR — begin the majority of their evaluations, opinion content is ambient noise. Original data is what gets cited, quoted, and retrieved.

This proposal recommends a twelve-month original research program built on the playbook HubSpot, Gong, and Drift used to convert proprietary data and category-defining surveys into inbound pipeline, analyst attention, and brand premium. Zenskar is unusually well positioned to run both halves of that playbook simultaneously:

- **The Drift play (category definition).** Zenskar has coined a category label — Zero-Touch Finance — but has not yet published the evidence that the category is real. A survey-based flagship study can do for Zero-Touch Finance what the State of Conversational Marketing did for Drift.
- **The Gong play (proprietary telemetry).** Zenskar's platform observes contract structures, usage events, billing runs, amendments, collections timing, and close cycles across a fast-growing customer base. That operating telemetry is a nascent Gong asset: data about how modern B2B monetization actually works that no analyst firm and no competitor can replicate.

We propose five sequenced research studies. The recommended first move is The State of AI-Era Billing — a 150-respondent survey of finance leaders at growth-stage B2B companies, validated synthetically before fielding, published under the Zero-Touch Finance thesis. It is the fastest study to execute, requires no customer-data governance work, and directly addresses the single highest-priority gap identified in the July 2026 LLM Training Data Audit of the billing cohort: original research is the one missing layer between Zenskar's current citation status of "Possible" and the "Likely Cited" status the incumbents enjoy.

There is also a clock running. Stripe's acquisition of Metronome (closed January 2026) has vacated the "independent usage-based billing authority" citation slot — including the annual survey franchise Metronome built. That slot will remain open for roughly one model-training cycle before Orb or another challenger claims it. The program below is designed to claim it first.

Situation Analysis: Zenskar in July 2026

The company

Zenskar is an AI-native billing and revenue automation platform founded in 2022 in New York by Apurv Bansal and Saurabh Agrawal. The platform models contract data as objects on a graph — covering usage tiers, prepaid credits, amendments, minimum commitments, customer hierarchies, and multi-entity, multi-currency structures — and pairs that structural flexibility with deterministic calculation and purpose-built AI agents across billing, revenue recognition, collections, and analytics. The 2026 product layer adds an Agents Marketplace, a Slack agent, and an MCP server connecting finance workflows to tools like Claude and ChatGPT. The positioning thesis, in Zenskar's own words, is Zero-Touch Finance: agents execute, humans supervise.

In April 2026 Zenskar closed a $15M Series A led by Susquehanna Venture Capital, Bessemer Venture Partners, Shine Capital, and Rho, bringing total funding to $21.5M, and reported 5x revenue growth in the trailing year. Customer proof points include Sardine (which replaced four years of in-house billing infrastructure), Pontera, Vertice, Thriva, Yembo, Posh, and 100ms. The ICP — heads of finance and controllers at usage-heavy SaaS, fintech, and AI companies in roughly the $5M–$100M ARR band — sits squarely in the demographic that now begins most vendor research through AI tools.

The competitive citation problem

Zenskar competes in a category whose top-of-funnel is consolidating around a handful of AI citation slots. The July 2026 LLM Training Data Audit of the nine-company billing and monetization cohort found that LLMs typically name only three to four brands per response, and the incumbent slots for subscription-billing queries are effectively held by Zuora, Chargebee, Stripe, and Recurly — three of them 2025 Gartner Magic Quadrant Leaders. Challengers compete for at most one rotating slot.

Zenskar's audit profile is the classic well-run challenger signature:

Dimension	Zenskar (Jul 2026)	What it means
Training Data Quality	**5 / 10**	Mid-cohort; strong ungated content and fresh funding citations, but no authority tier
Citation likelihood	**“Possible”**	Earns challenger-slot mentions; holds no incumbent slot
Analyst coverage	**None**	Gartner/Forrester inclusion is an 18+ month path; not a near-term lever
Peer reviews	**Thin–Moderate**	G2 reviews in the dozens, growing, below the ~50-review signal threshold
Original research	**None**	The audit's P0 remediation: the one missing layer between “Possible” and “Likely Cited”

The audit's cross-cohort finding is unambiguous: original primary research is the sharpest separator of citation authority in this category. Recurly converts a single asset class — ten years of State of Subscriptions data drawn from 76M subscribers and 2,200 merchants — into press-tier strength that its PE-quieted funding profile can no longer supply. Cohort members without original research show visibly thinner third-party citation trails. Zenskar has built the distribution machine; it has not yet built the thing that machine is best at distributing.

The window

Two structural events make the next two to three quarters unusually favorable:

- **The Metronome vacancy.** Stripe's ~$1B acquisition of Metronome, completed in January 2026, extinguished the category's independent usage-billing authority — and orphaned its State of Usage-Based Pricing survey franchise. The audit estimates the vacated slot stays open for roughly one model-training cycle. Whoever publishes the authoritative usage-pricing dataset next inherits the citation trail.
- **The AI-monetization narrative.** The dominant technology story of 2026 — how AI products are priced, metered, and monetized as token economics shift under every vendor — has no neutral data authority. Zenskar's ICP lives this problem daily. A company whose platform natively handles AI-token pricing volatility is the natural publisher of the data that defines it.

Set against these is a real constraint: Zenskar's Series A war chest and ~100-person team mean research investments must be sequenced, de-risked before fielding, and engineered so each study performs multiple jobs. The program below is designed around exactly those constraints.

The Playbook: Why Original Research Outperforms Content

The framework behind this proposal is documented in DevelopmentCorporate's analysis “They Published the Data. Then the Leads Came to Them,” which dissects how three enterprise software companies converted original research into category ownership at early stages:

- **HubSpot** launched the annual State of Inbound around 2010, when “inbound marketing” had no vocabulary and no analyst coverage. The report's statistics were each a stealth argument for the category thesis; over 5M downloads later, the research program was a core pillar of a ~$27B IPO brand.
- **Gong** recognized its call-recording telemetry as the most powerful content asset in its category. From ~2018 it published findings no competitor could replicate — because no competitor had the data. Within two years Gong was the most-cited company in sales productivity research, a material contributor to its $7.25B Series D valuation.
- **Drift** coined “conversational marketing,” then published the survey evidence that the category was real. Within twelve months the label was appearing in Gartner Market Guides and Forrester Wave reports — earned with zero incremental ad spend.

Three mechanisms make this playbook structurally superior to conventional content marketing, and all three have intensified since these case studies ran:

- **The backlink dynamic.** Journalists, analysts, and bloggers need primary sources. When your study is the origin of the statistic, every citation is an inbound link.
- **The GEO dynamic.** AI engines — the channel where Zenskar's buyers now begin their evaluations — preferentially cite primary sources when synthesizing answers. A statistic attributed to Zenskar in a well-structured, ungated report will be retrieved for years. This mechanism is the direct answer to the Training Data Audit's findings.
- **The category definition dynamic.** Publishing data about a problem is an implicit claim of authority over that problem. The company whose data defines Zero-Touch Finance will be perceived as its natural implementation.

Each study in this program is engineered to perform five simultaneous jobs: (1) a flagship report that anchors the launch; (2) derivative content assets — ungated summaries, LinkedIn packs, charts — that drive SEO and GEO; (3) press hooks built on counterintuitive, quotable findings; (4) sales enablement material that opens discovery calls and arms battlecards; and (5) category-definition claims that accumulate into analyst and investor perception. One study, five jobs.

Zenskar's Two Plays: Drift Now, Gong Compounding

Most companies can run only one half of this playbook. Zenskar can run both, on different clocks.

The Drift play — available immediately

Zenskar has already done the hard part of the Drift playbook: it has coined and consistently deployed a category label, Zero-Touch Finance, and articulated the thesis behind it (legacy systems were built for a simpler world; bolting AI onto broken foundations preserves their limitations). What is missing is the evidence layer — survey data quantifying how finance teams actually experience the problem: hours lost to billing operations, close-cycle length, revenue leakage from manual workarounds, the gap between AI ambition and AI adoption in the finance function. A co-sponsored survey study can be fielded within a single quarter and does not depend on customer-data governance.

The Gong play — an emerging proprietary asset

Zenskar's platform observes something no survey can: the actual anatomy of modern B2B monetization. Contract structures and their amendments, usage-event volumes, billing-run accuracy, credit and rollover behavior, collections timing, and revenue-recognition cycle times — across a customer base concentrated in exactly the segments (usage-heavy SaaS, fintech, AI) the market most wants data about. At today's scale this telemetry supports directional, anonymized benchmark findings; as the customer base compounds, it becomes a durable annual franchise that competitors cannot replicate. The honest sequencing is to lead with survey-based studies now and introduce platform-telemetry studies as the dataset deepens — with aggregation and anonymization standards established from day one.

One stakeholder-risk note, addressed in the study designs below: Zenskar's “scale operations, not headcount” narrative can read as a job-elimination story to the very finance professionals whose survey participation, social amplification, and buying committees the program depends on. Every study in this program frames automation as the elevation of finance work — grunt work to strategic work — and includes findings that quantify what finance teams gain, not merely what companies save. This mirrors the mitigation discipline we apply to workforce-sensitive data in other engagements.

The Five-Study Research Program

The five studies below are sequenced across twelve months. Each is designed around a category thesis rather than a product feature, engineered for quotable and counterintuitive findings, validated synthetically before real-respondent fielding, and built to perform the five jobs described above. All headline statistics shown are hypotheses to be validated — illustrative of the finding architecture, not established facts.

Study 1 — The State of AI-Era Billing

The anchor study — a Drift-style category-definition survey. Recommended first move.

Design

A 150-respondent survey of finance leaders (VP Finance, Controller, Head of RevOps) at B2B companies between $5M and $100M ARR, screened for usage-based or hybrid pricing. Twelve to fifteen quantitative questions covering billing-operations hours, close-cycle length, revenue leakage incidents, workaround prevalence, AI-tool adoption in the finance function, and the trust barriers holding it back. Co-sponsorship with one or two non-competing companies sharing the ICP (e.g., a CFO community, an accounting-automation adjacency) extends reach and splits panel costs, following the Drift/Salesforce model.

Hypothesized headline findings (to be validated)

- A majority of growth-stage finance teams still close the month in spreadsheets despite owning modern billing tools — quantifying the “bolt-on AI” thesis.
- Finance teams at usage-priced companies spend a materially larger share of hours on billing operations than subscription-only peers.
- A striking share of finance leaders report at least one revenue-leakage incident in the past year traced to a contract amendment their systems could not model.
- Finance trails every other function in AI adoption — and names auditability, not capability, as the blocker.

Why it leads

It executes fastest (one quarter, no data-governance dependencies), it directly discharges the Training Data Audit's P0, and it plants the Zero-Touch Finance category flag with evidence before any competitor publishes in the space. Every subsequent study inherits its distribution list and press relationships.

The five jobs

Flagship gated report (lead capture) · ungated findings summary (SEO/GEO) · press brief keyed to the AI-in-finance narrative · discovery-call statistics for sales · the founding evidence document for the Zero-Touch Finance category claim.

Study 2 — The Usage-Based Pricing Benchmark

Claiming the Metronome-vacated slot before Orb consolidates it.

Design

A hybrid study: a focused survey wave (~120 respondents, weighted toward AI and infrastructure companies) on pricing-model composition, metering pain, and repricing frequency, combined with anonymized directional telemetry from Zenskar's platform on usage-event patterns, credit behavior, and pricing-model mix. Published as an annual franchise — the successor dataset to the orphaned State of Usage-Based Pricing.

Hypothesized headline findings (to be validated)

- AI companies reprice materially more often than traditional SaaS — each model release is a billing event.
- Prepaid credits with expiry/rollover are becoming the dominant usage-monetization structure in the AI segment.
- A significant share of usage-priced companies cannot invoice a mid-term contract amendment without engineering involvement.

Why it matters

The audit identifies the vacated independent-usage-billing citation slot as Zenskar's single best near-term opening, open for roughly one training cycle. An annual benchmark franchise is the strongest possible claim on it — and directly feeds the comparison-content program Zenskar already runs.

The five jobs

Annual flagship benchmark · quarterly ungated data drops between editions · press hooks tied to AI-pricing news cycles · battlecard data for Orb/Chargebee matchups · the “independent usage-billing authority” positioning claim.

Study 3 — The Revenue Leakage Index

The first pure Gong-play study, built on platform telemetry.

Design

An anonymized, aggregated analysis of billing accuracy across Zenskar's customer base: leakage recovered at migration, error classes by pricing-model complexity, amendment-driven misbilling, and collections acceleration post-automation. Zenskar already asserts a 2–3% revenue uplift from error-free billing on its website; this study converts an unowned marketing claim into an owned, methodologically documented statistic — the difference between an assertion and a citation.

Hypothesized headline findings (to be validated)

- Companies migrating from spreadsheet-based or in-house billing recover a quantifiable percentage of previously leaked revenue in the first two quarters.

- Leakage risk scales super-linearly with contract complexity — amendments and minimum commits, not base subscriptions, are where revenue disappears.
- Collections timing improves by a measurable number of days once invoicing errors are eliminated.

Data governance

Requires customer consent architecture, aggregation thresholds (no finding attributable to fewer than N customers), and legal review — which is why it is sequenced third, after the consent and anonymization framework is built during Studies 1–2.

The five jobs

Flagship index report · a leakage self-assessment tool (interactive derivative asset) · CFO-press hooks (“the cost of billing you can't see”) · ROI evidence for sales proposals · the quantified backbone of the Zero-Touch Finance economic argument.

Study 4 — The Close Velocity Report

Benchmarking the month-end close in the automation era.

Design

A survey-plus-telemetry study on close-cycle length, ASC 606 automation rates, audit-preparation burden, and the close-time delta between automated and manual revenue-recognition workflows. Anchored by customer-verified outcomes Zenskar has already published (e.g., Vertice's 70% faster close) expanded into a population-level benchmark.

Hypothesized headline findings (to be validated)

- The median growth-stage B2B company closes in a number of days that has not improved in five years, despite tooling investment.
- Companies with automated rev-rec close in a fraction of the median — the “zero-day close” cohort is real and measurable.
- Audit preparation consumes a quantifiable multiple of hours at companies running spreadsheet-based ASC 606.

Why it matters

Close velocity is the metric CFOs already report to boards — it makes the program legible to the economic buyer, complements the controller-level appeal of Studies 1–3, and generates the finance-conference speaking material (Controllers Council, SaaSOps, CFO forums) that substitutes for absent analyst coverage.

The five jobs

Flagship benchmark · a close-velocity calculator · trade-press hooks in the accounting/CFO media lane · board-level ROI slides for sales · the operational proof pillar of Zero-Touch Finance.

Study 5 — The Finance Team of 2030

The category-institution study: what finance work becomes when agents execute.

Design

A forward-looking survey of 150+ finance leaders on role evolution, skills demand, agent oversight models, and headcount-vs-scope trajectories — deliberately framed around the elevation of finance work, not its elimination. Fielded to coincide with the second edition of Study 1, converting the program from a series of campaigns into an annual research institution.

Hypothesized headline findings (to be validated)

- Finance leaders expect a majority of transactional finance tasks to be agent-executed within a defined horizon — while strategic finance headcount grows.
- “Agent supervisor” emerges as a named finance role; a measurable share of teams are already writing it into job descriptions.
- Revenue-per-finance-FTE at automation-mature companies runs a quantifiable multiple above peers.

Stakeholder-risk mitigation

This is the study where the labor-sensitivity framing matters most. The instrument is designed with practitioner reviewers, findings are framed around career trajectory and skill premium, and the report includes a practitioner-facing chapter — ensuring the people the data describes become its amplifiers rather than its critics.

The five jobs

Flagship futures report · recruiting and employer-brand assets (a secondary payoff for Zenskar's own hiring) · national-press hooks on the future-of-work narrative · executive-conversation openers for founder-level selling · the long-horizon claim on the Zero-Touch Finance category.

Twelve-Month Program Calendar

Studies are sequenced so each launch inherits the audience, press relationships, and data infrastructure of the last, and so the platform-telemetry governance work (consent, aggregation, anonymization) is complete before the first Gong-play study fields.

Months	Milestone	Program activity
1–3	Study 1 fielded & launched	Category thesis workshop; instrument design; synthetic validation pass; real-respondent fielding; flagship launch with full five-job asset system
3–5	Study 2 fielded & launched	Usage-pricing survey wave; telemetry governance framework built in parallel; benchmark launch targeting the vacated usage-billing slot
5–7	Study 3 launched	Revenue Leakage Index from anonymized platform telemetry; interactive self-assessment derivative asset
7–9	Study 4 launched	Close Velocity Report; CFO-media and finance-conference push; board-level sales assets
9–12	Study 5 + annual cadence	Finance Team of 2030 futures study; planning for Year 2 editions of Studies 1 and 2 — converting campaigns into category institutions
Ongoing	Distribution & measurement	Monthly LLM citation-slot tracking (prompt battery across ChatGPT, Claude, Gemini, Perplexity); G2 review-velocity program toward the 50-review threshold; analyst-briefing groundwork

Measurement is built in from month one: the same live prompt-testing protocol used in the July 2026 Training Data Audit becomes the program's scoreboard, tracking Zenskar's movement from “Possible” toward “Likely Cited” across category-level and comparison prompts.

Methodology: Synthetic-First De-Risking

The most common reason Series A companies never execute a research program is fielding risk: a 150-respondent ICP-targeted panel costs roughly $3,500–$7,000 in panel fees, and if the findings come back flat, the investment — and the launch window — is wasted.

Every survey study in this program runs the Sandwich Method: before a dollar of panel spend, we generate 150–250 AI-synthetic respondents built on detailed ICP persona profiles and simulate the instrument. The synthetic pass answers three questions no other method can answer cheaply:

- Whether each question measures what we think it measures, or is ambiguous in ways that produce uninterpretable distributions;
- Whether the hypotheses will produce findings interesting and counterintuitive enough to publish — or so obvious they are unpublishable;
- Which specific statistics are most likely to generate press coverage and social sharing — the hooks.

The discipline is absolute: synthetic data is a design tool, never a data source. Published findings come exclusively from real respondents recruited through quality B2B panels (Respondent.io, Prolific, Cint) and, for telemetry studies, from consented, anonymized, threshold-aggregated platform data with documented methodology. For a company whose product thesis is “agents execute, humans supervise,” the research program's own methodology — AI accelerates the design, humans validate the evidence — is itself an on-brand story.

Commercial Terms

DevelopmentCorporate proposes three engagement options. All options use fixed-fee, milestone-based terms; each milestone is funded in advance before work on that milestone begins, and deliverables submitted for review are actioned within three business days. Panel fees for real-respondent fielding are passed through at cost. Pricing ranges below are indicative and will be finalized in the engagement agreement.

Engagement options

Option	Scope	Investment
Option A — Pilot Study (Study 1 only)	The State of AI-Era Billing end-to-end: thesis workshop, instrument design, synthetic validation, real-respondent fielding, flagship report, and the full five-job asset system. Approximately one quarter.	**$XX,XXX**
Option B — Year-One Program (Studies 1–3)	The anchor survey, the Usage-Based Pricing Benchmark, and the Revenue Leakage Index, including the telemetry governance framework and monthly LLM citation-slot tracking. Months 1–7.	**$XX,XXX–$XXX,XXX**
Option C — Annual Research Retainer (full five-study program)	All five studies, the twelve-month calendar, ongoing distribution support, citation-slot measurement, and Year-2 annual-edition planning.	**$XXX,XXX**

Milestone structure (illustrative, per study)

Milestone	Deliverable	Payment
1. Design & synthetic validation	Category thesis brief, survey instrument, synthetic-panel findings memo with identified hooks	30%
2. Fielding & analysis	Real-respondent dataset, cleaned and analyzed; findings review session	40%
3. Publication system	Flagship report, ungated summary, press brief, LinkedIn content pack, sales enablement extract	30%

To proceed, select an engagement option and we will issue an engagement agreement with the finalized milestone schedule. Work begins upon funding of Milestone 1.

About DevelopmentCorporate

DevelopmentCorporate LLC is an M&A advisory and strategic consulting firm specializing in early-stage SaaS companies. Managing Director John Mecke brings 30+ years of enterprise software experience, including executive roles at KnowledgeWare and Sterling Software, where he led over $300M in acquisitions. The firm's practice areas include competitive intelligence and battlecards, win/loss analysis, product-market-fit studies, LLM visibility planning, and research-driven demand generation programs built on the Evidence-First Research methodology.

The framework applied in this proposal is documented publicly at developmentcorporate.com in “They Published the Data. Then the Leads Came to Them” and the Sandwich Method series on hybrid synthetic-plus-human research.

> *Original data is the scarcest commodity on the internet. The companies that produce it do not chase traffic. Traffic comes to find them.*

Contact: John Mecke · Managing Director · DevelopmentCorporate LLC · developmentcorporate.com