

AI Consulting Business Plan

Mid-Market, Digital-First, Industry-Agnostic

Designed for a solo operator targeting \$500k-\$1M ARR
(leveraging a proven forecasting + segmentation case study)

Heaviside Digital (Cincinnati-based; national delivery)

January 29, 2026

Executive summary (no fluff)

- Positioning: “90-day AI implementation for mid-market teams that need production outcomes, not pilots.”
- Target buyer: COO/VP Ops, CTO/Head of Data, CFO (budget holder) at \$10M-\$250M revenue companies; AI-curious but stuck at pilots or ad-hoc tool use.
- Offer: productized engagement ladder (Assessment -> Pilot -> Production Launch -> Ops Retainer) with fixed scope, fixed timelines, and clear acceptance criteria.
- Go-to-market: 100% digital (LinkedIn + targeted outbound + SEO); no dependence on local networking or Cincinnati client base.
- Economics: You reach \$500k-\$1M ARR with 6-10 projects/year (\$50k-\$120k) plus 2-5 retainers (\$6k-\$15k/mo).
- Proof: your forecasting/segmentation case study is the anchor; additional credibility comes from widely cited industry outcomes and vendor case studies (see the Vertical Engagement Library).

Market signals to cite in sales conversations:

- AI is already used in at least one business function by a large share of organizations (McKinsey survey). [S1]
- Middle market leadership is actively exploring AI for customer insights, risk/fraud, efficiency, and productivity (RSM). [S2–S3]

What AI consulting work looks like (in practice)

Most paid engagements fall into four repeatable patterns. You want to sell patterns, not “custom AI.”

Engagement	Typical timeline	Typical fee (solo)	Deliverables clients expect
AI Opportunity Assessment	2-3 weeks	\$12k-\$25k	Use-case shortlist (3-5), data readiness scorecard, ROI model, 90-day roadmap, risk & governance notes
Rapid Pilot / MVP	6-8 weeks	\$35k-\$65k	Working prototype, baseline vs model metrics, deployment plan, operator runbook, user training, go/no-go decision
Production Launch (single use case)	8-12 weeks	\$65k-\$140k	Production system, integrations, monitoring + alerting, security review artifacts, documentation, handoff
AI Ops Retainer	monthly	\$6k-\$15k/mo	Model monitoring, prompt/model drift checks, quarterly roadmap refresh, backlog delivery (small improvements)

A practical 90-day delivery shape (for ‘Production Launch’):

- Week 1-2: discovery + data audit + success metrics + architecture decision + SOW lock.
- Week 3-6: build (data pipeline + model/prompt + evaluation) + weekly demos.
- Week 7-9: integration (APIs, security, logging) + acceptance tests.
- Week 10-12: rollout (training, monitoring dashboards, handover) + retainer conversion conversation.

Why clients pay:

- They need a production outcome with guardrails (security, privacy, auditability) and internal teams are stretched.
- They want to avoid ‘pilot purgatory’—something works in a notebook but never ships.
- They’ll pay more for speed + clarity: fixed deliverables, fixed timeline, and measurable acceptance criteria.

Who pays (ICP) and what problems are worth \$50k-\$150k

- Company size: \$10M-\$250M revenue; 50-2,000 employees; data exists but is messy; leadership wants measurable ROI in one quarter.
- Budget trigger: measurable savings/revenue within 6-12 months OR compliance/risk reduction that de-risks the business.
- Buying motion: 1-3 stakeholders; avoids long procurement; can sign a fixed-scope project with a solo specialist.

Common mid-market buying triggers (cross-industry):

- Forecasting is unreliable (inventory, staffing, revenue, demand).
- Operations are bottlenecked by unstructured docs (claims, contracts, tickets, emails).
- Support costs rising; NPS slipping; handle time too high.
- Fraud/chargebacks or compliance reviews growing.
- Leadership wants a knowledge assistant but needs it grounded in internal policy and data.

Middle market context you can cite:

- RSM reports on how middle market leaders are using AI and where they see impact. [S2-S3]
- McKinsey's annual AI survey reports broad adoption and where AI is being used inside companies. [S1]

How it's priced (models, ranges, tradeoffs)

Pricing model	When it works	Typical range (solo)	Tradeoffs
Fixed-fee project	Defined deliverables + timeline	\$35k-\$140k	Best margins; requires strong scope control
Monthly retainer	Ongoing optimization & monitoring	\$6k-\$15k/mo	Predictable; can become 'support desk' unless scoped
Day rate	Workshops / discovery sprints	\$1.5k-\$3k/day	Easy to sell; lower leverage than fixed-fee
Value-based (shared upside)	Measurable ROI use cases (forecasting, fraud)	Base + success fee	High upside; requires clean measurement and client trust

Fee-setting benchmarks to cite:

- Consulting Success' 2023 fee study (updated in 2025) shows value-based pricers skew toward higher project values; hourly users cluster lower. [S4, S27]

Vertical engagement library: what's being sold (and cited examples)

Use this section as: (1) inspiration for your offer stack, (2) proof points in sales calls, and (3) the basis for vertical-specific landing pages. Each vertical includes: top engagements, typical deliverables, and citeable outcomes.

Customer Service / Contact Center

Most common paid engagements:

- GenAI agent for Tier-1 support (deflection + routing + agent assist).
- Call reason prediction and smart routing.
- Knowledge base cleanup and governance (so answers are trustworthy).

Cited example	What was built	Reported outcomes (high level)	Source
Klarna AI assistant	GenAI assistant for customer service chats.	2.3M conversations; ~2/3 of chats; ~700 FTE equivalent; 25% fewer repeat inquiries; <2 min resolution vs 11 min.	[S5]
Verizon GenAI call routing	Predict call reasons and route to best agents; speed store visits.	Goal: predict ~80% of call reasons; target: prevent 100,000 customers leaving in 2024; ~7 min shorter store visits.	[S26]

Financial Services (Banking / Wealth / Payments)

Most common paid engagements:

- Internal knowledge assistant for advisors/analysts (RAG over policy + research).
- Fraud / AML triage copilots (case summarization, evidence gathering).
- Customer service automation with escalations and audit trails.

Cited example	What was built	Reported outcomes (high level)	Source
Morgan Stanley Assistant	GPT-4 knowledge assistant for advisors.	Over 98% of advisor teams actively use it.	[S6]
Mastercard fraud detection on AWS	AI/ML fraud detection at massive scale.	Detect 3x fraud; reduce false positives 10x; billions in merchant savings (as reported by AWS).	[S7]

Insurance

Most common paid engagements:

- Claims & document triage automation (OCR + classification + GenAI summaries).
- Fraud detection / anomaly detection (transaction + behavioral + graph signals).
- Underwriting copilots (policy Q&A grounded in internal rules; risk scoring assist).

Cited example	What was built	Reported outcomes (high level)	Source
Shift Technology	Claims/fraud documentation processing with AI vision + GenAI.	Reduced weeks of sorting/classification/decisions to days.	[S8]
legal-i	Medical insurance claims processing support with GenAI.	Up to 80% faster; up to 4x accuracy; payouts optimized up to 11%.	[S9]
Amazon Fraud Detector (Wuzzon)	Fraud detection platform integrated into marketing attribution.	False positives down up to 43%; true positive rate up 11-14% (varies by source).	[S10]

Healthcare (Providers / Payers)

Most common paid engagements:

- Ambient documentation and note drafting (speech -> structured notes -> EHR).
- Capacity / flow optimization (ED wait, scheduling, triage).
- Payer automation: prior auth, policy interpretation, internal process docs.

Cited example	What was built	Reported outcomes (high level)	Source
Hero AI	ER workflow optimization.	55% reduction in patient wait times; +200 hours ER capacity.	[S11]
healow Sunoh.ai	Ambient AI scribe for clinical notes.	Nearly 50% reduction in documentation time; up to ~2 hours/day saved.	[S12]
Klinikum Landsberg	Real-time transcription + structured documentation.	Documentation effort reduced by >90%; up to ~2 hours/day freed.	[S13]
CareSource	Payer automation and copilots.	Process reduced from ~2 months to ~2 weeks; >\$125k saved by one automation initiative.	[S14]

Retail / E-commerce

Most common paid engagements:

- Demand forecasting & replenishment improvements (time-series + promo effects).
- Search/recommendation/personalization (next best product/content).
- Customer support + returns automation (ticket summarization, routing).

Cited example	What was built	Reported outcomes (high level)	Source
Central Group	AI search/promo/stock assistant.	Search time -94% (32 min -> 1-2 min); conversion +10%.	[S15]
Amazon Personalize (Viewlift)	Recommendations and personalization.	+24% clicks on recommended videos tray (vs curated/auto).	[S16]
Albert Heijn forecasting	ML demand forecasting modernization.	Reported +5% prediction accuracy.	[S17]

Manufacturing / Industrial

Most common paid engagements:

- Predictive maintenance (sensor time-series; failure prediction; alerts).
- Demand/revenue forecasting and product/customer segmentation (your case study maps here).
- Quality inspection analytics (vision models; defect detection) where data exists.

Cited example	What was built	Reported outcomes (high level)	Source
Husky Technologies	Remote monitoring/analytics to preempt production issues.	Estimated \$4k-\$6k savings per intervention.	[S18]
Fincantieri	Predictive maintenance for turbogenerators.	Case study emphasizes less downtime and improved maintenance decisions.	[S19]

Logistics / Transportation

Most common paid engagements:

- Route optimization and dispatch assistance (constraints + ML forecasting).
- Warehouse picking optimization (slotting, demand heatmaps).
- Customer ETA prediction and exception management (delay causes + next best action).

Cited example	What was built	Reported outcomes (high level)	Source
UPS ORION	Route optimization rollout program.	Expected annual reduction: 100M miles; 10M gallons fuel; 100k metric tons GHG avoided (per BSR case study).	[S20]
DHL route optimization	Route/stop optimization to reduce unnecessary distance, time, cost, emissions.	Describes measurable reductions in unnecessary kilometers and operating time.	[S21]

Professional Services / Legal / Compliance

Most common paid engagements:

- Contract review and clause extraction; risk flagging; drafting assistance.
- Research copilots (grounded in internal memos/templates/precedent).
- Compliance evidence automation (policy mapping; controls narratives).

Cited example	What was built	Reported outcomes (high level)	Source
JPMorgan COIN	ML contract intelligence for loan agreements.	Reported savings: ~360,000 hours/year.	[S22]
Fidal IA	Legal research assistant.	Legal research time reduced ~30%.	[S23]
DraftWise	Legal drafting/search tools.	Developer productivity +60% (reported).	[S24]

Travel / Hospitality (Pricing)

Most common paid engagements:

- Common patterns: knowledge assistant, document workflow automation, forecasting, risk detection.

Cited example	What was built	Reported outcomes (high level)	Source
airBaltic dynamic pricing	AI-driven seat pricing automation.	6% revenue increase; results achieved in ~6 months.	[S25]

Your productized offers (recommended)

These are designed to (a) fit a solo schedule, (b) reduce scope creep, and (c) ladder clients into a retainer.

Offer	Who it's for	Timeline	Price	Success criteria (examples)
AI Opportunity Assessment	Exec team wants a short list + ROI + next step	2-3 weeks	\$15k-\$25k	Top 3 use cases, data readiness score, ROI model, delivery plan
Pilot-to-Production Sprint	Team has a use case but needs a working MVP	6-8 weeks	\$40k-\$65k	MVP deployed to staging; evaluation report; go/no-go to production
90-Day Production Launch	One use case ready for production	8-12 weeks	\$75k-\$140k	Live in production with monitoring; adoption metric agreed; handoff complete
AI Ops Retainer (Optimization)	Client wants ongoing value + safety	monthly	\$6k-\$15k/mo	Monthly KPI review; drift checks; backlog delivered; quarterly roadmap refresh

How your proprietary case study fits:

- Your forecasting + segmentation project becomes the blueprint for (1) assessments, (2) pilots, and (3) production launches in any vertical with historical transaction/sales/usage data.
- Build a de-identified one-pager and a longer technical appendix (data volume, features used, error metrics, business lift). Use it everywhere.

12-24 month revenue model (solo) with clear assumptions

These are planning models, not promises. Adjust conversion rates and deal sizes as you get real pipeline data.

Assumptions:

- Average project length: 8-10 weeks; you can run 1 large or 2 small projects concurrently if scoped tightly.
- Retainer conversion: 20-35% of projects convert into a retainer within 60 days of launch.
- Marketing time: 10-12 hours/week minimum (digital only).
- Close rate on qualified opportunities: 20-35% depending on proof and urgency.

Scenario	Projects/year	Avg project fee	Retainers	Avg retainer	ARR outcome
Conservative (hit ~\$500k)	6	\$60k	1-2	\$6k/mo	~\$504k-\$648k
Base (hit ~\$750k)	8	\$70k	2-3	\$8k/mo	~\$736k-\$1.0M
Aggressive (hit ~\$1M)	10	\$80k	3-5	\$10k-\$15k/mo	~\$1.0M-\$1.7M

Example 12-quarter ramp (base case):

Quarter	Projects closed	Avg fee	Active retainers (end of qtr)	Monthly retainer	Quarter revenue
Q1	1	\$45k	0	-	\$45k
Q2	1	\$55k	1	\$6k	\$73k
Q3	2	\$60k	1	\$8k	\$136k
Q4	2	\$65k	2	\$8k	\$178k
Q5	2	\$70k	2	\$10k	\$200k
Q6	2	\$75k	3	\$10k	\$240k
Q7	2	\$80k	3	\$12k	\$268k
Q8	2	\$85k	4	\$12k	\$314k

Marketing plans (digital-only; minimal networking time)

You're using your marketing skillset to sell AI implementation. The goal is a repeatable pipeline, not "brand building."

Plan A: LinkedIn-led (fastest to revenue)

- Daily: 1 post (300-700 words) + 10 thoughtful comments on posts by your target personas.
- Weekly: 1 asset promoted (case study snippet, checklist, or teardown).
- Paid: \$2k-\$5k/month to boost best-performing posts into the exact titles/company sizes you want.
- CTA: book a 20-min "AI Opportunity Triage" call (Calendly); qualify hard.

Plan B: Outbound-led (most controllable)

- Build a list of 300-600 targets/month (50-2,000 employees; ops-heavy; AI-curious).
- Send 50-75 highly targeted emails/week; personalize the first sentence only.
- Offer: fixed-fee 2-3 week Assessment; deliver a 1-page "Use case + ROI + 90-day plan."
- Follow-up sequence: 5 touches over 14 days; then recycle with a new insight 60 days later.

Plan C: SEO-led (compounding, slower)

- Publish 2 in-depth pages/month: one vertical page + one "how-to" (forecasting, document workflows, knowledge assistants).
- Build landing pages per vertical using the case studies in this doc, with clear 'what we ship in 90 days.'
- Capture leads with 1 lead magnet: "AI Implementation Readiness Checklist (mid-market)."

Weekly operating system (12-15 hours)

Task	Time	Output
LinkedIn content + engagement	5-6 hrs	5 posts + 50 targeted comments
Outbound prospecting + follow-ups	3-4 hrs	50-75 new contacts + 5-touch sequence
1 long-form asset or landing page	3-4 hrs	case study/vertical page/lead magnet update
Pipeline review + proposal writing	1-2 hrs	next steps, proposals, scheduling

Key risks, scalability limits, and mitigations

Risk	What it looks like	Mitigation (solo-friendly)
Scope creep	'Can you also...' requests expand effort	Fixed-scope packages; change-order policy; weekly acceptance checkpoints
Data reality gap	Data missing/dirty slows delivery	Data-readiness gate in Assessment; 'data sprint' line item; clear assumptions
Compliance/security stall	IT/security blocks deployment	Use RAG with citations; log everything; documented controls; early security review
Feast/famine pipeline	No leads when busy; busy when no leads	Non-negotiable 10 hrs/week marketing; systemize outreach; keep ads always-on
Solo capacity cap	Too many concurrent projects	Productize; cap WIP; use vetted contractors for narrow tasks (data eng, UI)

30 / 60 / 90-day action plan

Days 1-30 (assets + launch)

- Write your de-identified forecasting + segmentation case study (1 page + 6-slide mini deck).
- Create 3 landing pages: (1) Assessment, (2) 90-Day Production Launch, (3) 'Forecasting & Segmentation' flagship.
- Set up: CRM + Calendly + email sequences + basic analytics.
- Start daily LinkedIn posting + first outbound batch (50/week).

Days 31-60 (pipeline + first closes)

- Run 6-10 discovery calls; close 1 Assessment; deliver it fast and document everything.
- Turn assessment output into a 'Pilot-to-Production' proposal within 48 hours.
- Launch a small LinkedIn boost campaign to your best post (\$20-\$50/day).

Days 61-90 (ship + prove)

- Ship first Pilot or Production Launch with weekly demos and explicit acceptance criteria.
- Publish 2 more "external" case study teardowns (from this doc) to show cross-vertical breadth.
- Convert the first client into a retainer with a clear "optimization backlog" and monthly KPI review.

References (URLs)

- [S1] McKinsey, The State of AI (2025 PDF) — https://www.mckinsey.com/~media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai/2025/the-state-of-ai-how-organizations-are-rewiring-to-capture-value_final.pdf
- [S2] RSM Middle Market AI Survey 2025 — <https://rsmus.com/insights/services/digital-transformation/rsm-middle-market-ai-survey-2025.html>
- [S3] RSM: AI in the C-suite (middle market) use cases — <https://rsmus.com/insights/services/digital-transformation/ai-for-c-suite-leaders.html>
- [S4] Consulting Success: Consulting fees & value-based pricing stats (2023 study, updated 2025) — <https://www.consultingsuccess.com/consulting-fees>
- [S5] OpenAI customer story: Klarna AI assistant — <https://openai.com/index/klarna/>
- [S6] OpenAI customer story: Morgan Stanley — <https://openai.com/index/morgan-stanley/>
- [S7] AWS case study: Mastercard AI/ML fraud detection — <https://aws.amazon.com/solutions/case-studies/mastercard-ai-ml-testimonial/>
- [S8] Microsoft customer story: Shift Technology (claims docs weeks to days) — <https://www.microsoft.com/en/customers/story/23202-shift-technology-azure-ai-vision>
- [S9] Microsoft customer story: legal-i (insurance claims up to 80% faster) — <https://www.microsoft.com/en/customers/story/21393-legal-i-azure-open-ai-service>
- [S10] AWS Fraud Detector customers (Wuzzon; false positives -43%) — <https://aws.amazon.com/fraud-detector/customers/>
- [S11] Microsoft customer story: Hero AI (patient wait times -55%) — <https://www.microsoft.com/en/customers/story/22803-hero-ai-azure-open-ai-service>
- [S12] Microsoft customer story: healow Sunoh.ai (documentation time ~-50%) — <https://www.microsoft.com/en/customers/story/24772-healow-azure>
- [S13] Microsoft customer story: Klinikum Landsberg (documentation effort >-90%) — <https://www.microsoft.com/en/customers/story/25038-klinikum-landsberg-am-lech-azure-openai>
- [S14] Microsoft customer story: CareSource (2 months -> 2 weeks; \$125k savings) — <https://www.microsoft.com/en/customers/story/25041-caresource-azure>
- [S15] Google Cloud customer story: Central Group (conversion +10%, search time -94%) — <https://cloud.google.com/customers/central-group>
- [S16] AWS Personalize customers (e.g., Viewlift +24% clicks) — <https://aws.amazon.com/personalize/customers/>
- [S17] Xebia customer story: Albert Heijn forecasting (+5% accuracy) — <https://xebia.com/customer-stories/pioneering-grocery-chain-improves-prediction-accuracy-by-5/>
- [S18] Microsoft customer story: Husky Technologies (save \$4k-\$6k per intervention) — <https://www.microsoft.com/en/customers/story/1596648613315266488-huskytechnologies-azure-iot>
- [S19] Microsoft customer story: Fincantieri predictive maintenance — <https://www.microsoft.com/en/customers/story/1655460606196709857-fincantieri-energy-azure-en-it>

[S20] BSR case study: UPS ORION route optimization (100M miles; 10M gallons expected) —
<https://www.bsr.org/en/case-studies/center-for-technology-and-sustainability-orion-technology-ups>

[S21] DHL: route optimization algorithm (Greenplan) —
<https://www.dhl.com/global-en/delivered/responsibility/route-optimization-algorithm.html>

[S22] ABA Journal (via Bloomberg): JPMorgan COIN saves 360,000 hours —
https://www.abajournal.com/news/article/jpmorgan_chase_uses_tech_to_save_360000_hours_of_annual_work_by_lawyers_and

[S23] Microsoft customer story: Fidal IA (legal research -30%) —
<https://www.microsoft.com/en/customers/story/25200-fidal-azure-open-ai>

[S24] Microsoft customer story: DraftWise (developer productivity +60%) —
<https://www.microsoft.com/en/customers/story/23918-draftwise-azure-ai-foundry>

[S25] PROS case study: airBaltic AI pricing (+6% revenue) —
<https://pros.com/learn/case-studies-testimonials/how-airbaltic-driving-seat-revenue-with-ai-powered-dynamic-ancillary-pricing/>

[S26] Reuters: Verizon uses GenAI to improve loyalty (call reasons ~80%; prevent churn) —
<https://www.reuters.com/technology/artificial-intelligence/verizon-uses-genai-improve-customer-loyalty-2024-06-18/>

[S27] Consulting Success: hourly vs value-based project value comparison (same study) —
<https://www.consultingsuccess.com/consulting-fees>