

ILLUSTRATIVE BUSINESS DESIGN CASE

Designing the Business Behind an AI Product

Asterlane AI, Inc.

A fictional founder-led AI company considering a product that helps B2B software teams complete enterprise security, procurement, and compliance questionnaires using a verified company knowledge base and source-linked AI answers.

Stage	Pre-product / design	Initial buyer	Revenue Operations / Security
Target account	B2B SaaS, 50-500 employees	Primary motion	Founder-led direct
Proposed model	Annual SaaS + onboarding	Case date	September 2026

Business Design

Fellstead · Operated by Solten Ventures LLC

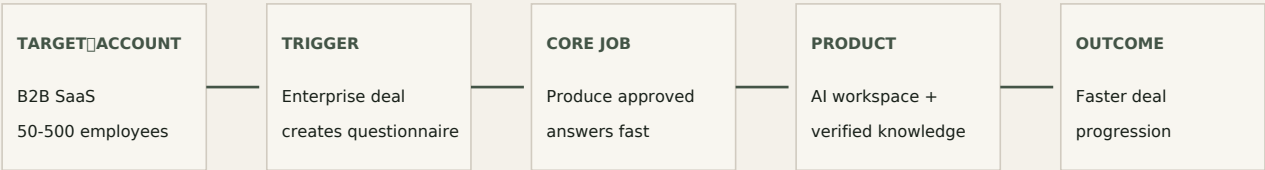
BUSINESS DESIGN QUESTION

The product idea is plausible. The business around it is not yet designed.

Asterlane's founders have a technically credible concept: use retrieval and generative AI to answer repetitive enterprise security, procurement, and compliance questionnaires from a controlled body of company-approved information. Early conversations suggest the problem is real. Several teams report spending substantial time locating prior answers, verifying policy language, routing questions to subject-matter experts, and formatting responses.

But a product concept is not yet a business. The founders have not resolved which buyer owns the pain, which account profile experiences it frequently enough to pay, where AI can operate autonomously versus requiring verification, how onboarding should work, what should be priced, how model cost affects margins, how the product enters an account, or which evidence would justify building beyond an initial wedge.

Proposed business architecture



Design principle: narrow the first business around one painful transaction, then expand from verified behavior rather than platform ambition.

Design objective: define a business that can be tested before committing a full engineering roadmap - customer, wedge, value, pricing, economics, go-to-market, delivery, operating model, metrics, and explicit expansion gates.

HOW TO READ THIS CASE

Business Design is a set of connected choices, not a feature list.

This case does not attempt to predict whether Asterlane will succeed. It demonstrates how Fellstead would turn a promising opportunity into a coherent business design that can be falsified, tested, and revised before substantial resources are committed.

Design layer	Core question
Customer	Who has the problem often enough, urgently enough, and visibly enough to buy?
Value	What measurable outcome makes the product worth adopting and renewing?
Offer	What exactly is sold, to whom, and with what boundaries?
Economics	Can pricing support AI usage, onboarding, support, and acquisition?
Go-to-market	How does the business find accounts at the moment the problem is active?
Operating model	What must humans do, what can AI do, and who owns exceptions?
Validation	Which evidence must be true before the business expands?

Illustrative case note

Asterlane AI is fictional. All metrics, interviews, pricing, economics, customer observations, and proposed decisions are simulated for demonstration. The purpose is to show Fellstead's Business Design method and work product, not to represent a client engagement or realized result.

OPPORTUNITY FRAME

Start with a transaction, not with 'AI for enterprise knowledge.'

The broad opportunity could be described as enterprise knowledge automation. That category is too large to guide an early business. The proposed wedge is narrower: helping B2B SaaS companies respond to customer security, procurement, and compliance questionnaires during enterprise sales cycles.

Dimension	Initial design choice	Reason
Trigger	A live enterprise deal generates a questionnaire or evidence request.	Pain is active, time-bound, and tied to revenue.
Primary user	Revenue operations / security / deal desk coordinator.	Owens coordination even when subject experts provide content.
Economic buyer	CRO, COO, or security leader depending on org design.	Can connect labor savings and deal velocity to budget.
Job to be done	Produce an accurate, approved response package quickly.	Specific outcome; easier to measure than generic knowledge search.
Boundary	Questionnaire response, evidence retrieval, approval workflow.	Avoids starting as a broad internal AI assistant.
Expansion option	Policies, customer trust center, RFPs, legal/procurement content.	Only after the wedge proves repeatable.

Design thesis

The business should enter through a high-friction, revenue-linked workflow where source accuracy and organizational memory matter. The first product should not promise autonomous enterprise knowledge. It should promise faster verified completion of a specific external request.

CUSTOMER DESIGN

The best first customer is not 'any company with questionnaires.'

Attribute	Target condition	Why it matters
Company	B2B SaaS / software infrastructure, 50-500 employees	Frequent enterprise sales, enough process pain, limited specialist staff.
Sales model	Meaningful enterprise or upper-mid-market motion	Questionnaires recur and are tied to important deals.
Volume	~20+ security/procurement packages per year	Enough repetition for annual software value.
Existing knowledge	Prior questionnaires, policies, security docs, product answers	Creates source corpus for controlled AI retrieval.
Current process	Spreadsheet/docs + Slack/email + expert routing	Visible coordination burden and searchable baseline.
Buyer condition	No fully staffed formal trust/RFP team	Pain is large relative to available people.
Avoid initially	Heavily regulated firms requiring bespoke legal review for most answers	Slower sales and higher liability before controls mature.

Qualification rule

Asterlane should qualify accounts by workflow frequency and coordination burden, not simply employee count or industry. A smaller company answering 60 questionnaires can be a better customer than a larger company answering five.

VALUE ARCHITECTURE

The value proposition must survive without the word AI.

AI is the mechanism. The customer value should be framed around business outcomes that can be observed before and after adoption.

Value layer	Proposed outcome	Evidence to collect
Speed	Reduce median completion cycle from days to hours for repeatable questions.	Start/finish timestamps; time-to-first-draft.
Expert load	Reduce repeated requests to security, legal, product, and engineering experts.	Number of routed questions; expert minutes per package.
Consistency	Reuse approved answers without copy/paste drift.	Answer reuse rate; approval changes; source coverage.
Deal support	Remove questionnaire response as a preventable source of deal delay.	Sales-stage aging attributable to response workflow.
Memory	Preserve why an answer was approved and when it must be revisited.	Source links; owner; validity date; change history.

Primary promise

For B2B software teams selling to enterprise customers, Asterlane turns approved company knowledge into source-linked questionnaire responses - reducing coordination and expert rework while preserving human control over what leaves the company.

PRODUCT BOUNDARY

Design the human-AI boundary before designing the roadmap.

Task	AI role	Human role	Release rule
Retrieve prior answer	Find semantically relevant approved content.	None if source is current.	Can automate retrieval.
Draft response	Compose answer from approved sources.	Review when confidence/source coverage below threshold.	Draft, not autonomous external send.
Evidence attachment	Suggest relevant policy / certification / document.	Confirm correct version and disclosure status.	Human confirmation required initially.
New / novel question	Generate candidate from source material.	Subject expert authors/approves canonical answer.	Must create new approved knowledge.
Sensitive / legal language	Surface prior approved language only.	Legal/security owner reviews.	No free-form autonomous answer.
Final submission	Assemble response package.	Named owner approves external release.	Human release always required in first design.

Why this matters

Asterlane should not make autonomy the product promise. The near-term business is stronger if trust comes from controlled sources, visible confidence, explicit ownership, and a faster approval process. Autonomy can expand only where observed error cost is low and evidence supports it.

OFFER DESIGN

Sell a controlled response workspace, not an AI seat.

Component	Included in initial offer
Knowledge setup	Import prior questionnaires, policies, standard documents, approved answers, ownership metadata.
Response workspace	Upload questionnaire; classify questions; retrieve sources; draft answers; track unresolved items.
Approval workflow	Route low-confidence or sensitive questions to named owners; preserve decision history.
Evidence library	Source-linked documents with owner, date, status, and refresh reminders.
Analytics	Cycle time, reuse rate, unanswered-question categories, expert routing, stale knowledge.
Onboarding	Structured first knowledge build + two live questionnaire workflows.
Not included initially	General-purpose internal search, autonomous outbound submission, full legal review, broad RFP automation.

Offer logic

The unit of value is not a user seat. It is a company workflow supported by shared approved knowledge. Seat-only pricing would underprice high-value shared usage and discourage cross-functional participation.

PRICING ARCHITECTURE

Price the workflow and account value, not token consumption.

Plan	Target account	Illustrative price	Design logic
Launch	50-150 employees / moderate volume	\$12k annual + \$2k onboarding	Low procurement friction; proves annual willingness to pay.
Growth	100-500 employees / frequent enterprise deals	\$18k-\$24k annual + \$3k onboarding	Core target; supports shared users and higher questionnaire volume.
Scale	High volume / multiple business units	\$36k+ annual + scoped onboarding	Only after controls, usage, and service burden are proven.

Pricing rules

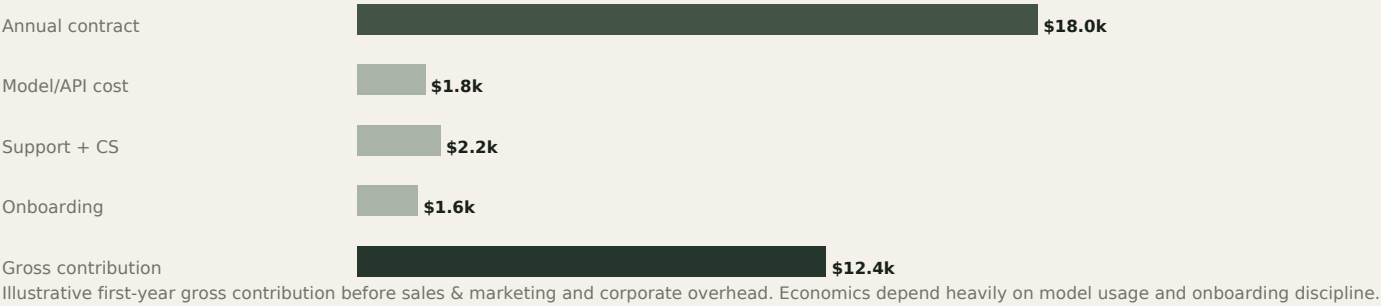
Do not expose model tokens as the primary pricing unit. Include a reasonable usage envelope and monitor cost internally. Do not use unlimited custom onboarding to close early accounts. Discount only against explicit learning value, contract term, or referenceability - not because pricing is undefined.

The initial design should test whether \$18k annual feels small relative to enterprise-deal coordination cost while remaining meaningful enough to support human onboarding and support.

UNIT ECONOMICS

AI margins are designed, not assumed.

Illustrative account economics at target pricing



Economic variable	Illustrative target	What must be controlled
Annual contract	\$18k median target	Avoid drifting to low-price seat SaaS.
Model / API cost	≤10% revenue	Prompt architecture, caching, retrieval, model routing.
Onboarding labor	≤\$1.6k first year	Standard import and canonical-answer workflow.
Support / CS	≤\$2.2k annual	Self-serve administration; bounded expert support.
Gross contribution	~69% before S&M; / overhead	Requires disciplined onboarding and model usage.
CAC payback target	<12 months once repeatable	Founder sales can mask true acquisition cost initially.

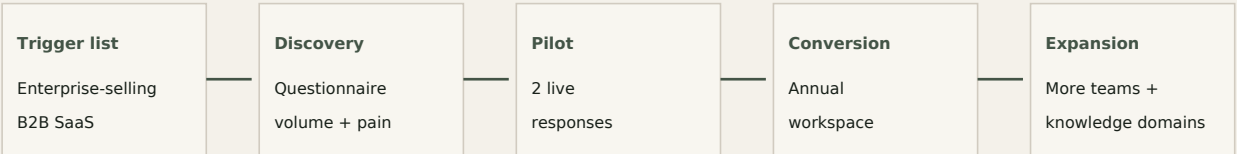
Critical economic risk

The product can look like high-margin software while quietly becoming a managed knowledge service. Every manual research, document-cleanup, policy-writing, or answer-authoring task must be tracked separately during pilots so software economics are not inferred from blended delivery.

GO-TO-MARKET DESIGN

Acquire customers when the pain is visible, not when they are browsing for AI tools.

Initial go-to-market sequence



Channel	Initial use	Reason
Founder-led outbound	Target companies visibly selling enterprise; message around questionnaire burden.	Fastest learning loop and precise qualification.
Security / compliance advisors	Referral when clients repeatedly assemble evidence and answers.	Advisor sees the workflow pain before software vendors do.
Fractional RevOps / deal desk	Referral for teams scaling enterprise sales without dedicated response ops.	Strong trigger visibility.
Content / case study	Show before/after workflow and knowledge-control model.	Proof of work; supports trust.
Product-led signup	Not primary at launch.	Problem requires setup, shared knowledge, and multi-stakeholder adoption.

Message

Do not lead with 'AI compliance automation.' Lead with the moment: 'Your enterprise buyer sent another 180-question security package. How many people will touch it before it goes back?'

SALES MOTION

The first sale should validate the workflow before it validates the platform.

Stage	Purpose	Evidence required to advance
Trigger discovery	Confirm frequency, ownership, cycle time, expert routing, and business impact.	At least 20 relevant packages/year or comparable recurring burden.
Workflow review	Map one recent questionnaire end to end.	Visible repeated answers and coordination handoffs.
Pilot	Run two live questionnaires with controlled knowledge base.	Measured cycle time, source coverage, review burden, answer reuse.
Business case	Compare labor/coordination reduction and deal-stage friction.	Buyer agrees outcome matters enough for annual budget.
Annual conversion	Standard workspace + onboarding + support boundary.	Named owner, approved sources, repeatable workflow.
Expansion	More domains, teams, or workflows.	Evidence of retention and usage before product breadth.

Pilot rule

A pilot is not free consulting. The goal is to test product behavior, not to manually complete questionnaires on the customer's behalf. Any manual work performed to make the pilot succeed must be logged as delivery cost and treated as a product-design signal.

OPERATING MODEL

The early company should be organized around learning loops, not departments.

Role / capability	Initial ownership	Core responsibility
CEO / founder	Commercial thesis + customer learning	ICP, pricing, sales, design decisions, expansion gates.
Technical founder	Product + AI system	Retrieval quality, model routing, security architecture, instrumentation.
Product / implementation lead	Customer setup + workflow design	Knowledge import, onboarding, support taxonomy, pilot learning.
Security advisor / contractor	Control design	Review security posture and disclosure boundaries; not daily delivery.
Design partner customers	Evidence source, not roadmap authority	Provide real workflows; do not convert every request into product.

Decision discipline

The founders should maintain an explicit design log: customer assumptions, pricing assumptions, AI-confidence thresholds, onboarding exceptions, and expansion gates. The most dangerous early-stage pattern is to encode one design partner's process into the product before deciding whether the behavior is segment-wide.

METRICS THAT MATTER

Instrument the business questions, not only product activity.

Metric	Question it answers
Questionnaire cycle time	Does Asterlane materially accelerate the job?
Draft acceptance rate	Are AI-generated answers useful before human review?
Source coverage	How much of the workflow can be grounded in approved knowledge?
Expert routing rate	Where is the knowledge base incomplete or risk threshold too low?
Canonical-answer creation rate	Is the system getting smarter through real work?
Knowledge staleness / refresh rate	Can approved memory remain trustworthy over time?
Active questionnaire accounts / month	Is usage tied to recurring workflow rather than novelty?
Pilot-to-annual conversion	Does workflow value translate to budget?
Gross contribution by account	Are AI and onboarding costs compatible with pricing?
Expansion by domain/team	Does broader platform value emerge from the wedge?

Do not optimize early for

Prompt counts, total generated words, raw user seats, or generic engagement. A product can produce large AI activity while failing to reduce the customer's actual coordination cost.

90-DAY VALIDATION PLAN

Make the business earn the right to be built.

The goal is to invalidate weak assumptions cheaply. Asterlane should not interpret interest, demo enthusiasm, or pilot participation as proof of a scalable business.

Days	Test	Action	Decision rule
0-30	Problem frequency	Interview 15-20 target accounts and reconstruct last questionnaire workflow.	Proceed if pain recurs frequently and coordination burden is material.
0-30	Buyer ownership	Test RevOps/security/deal-desk ownership and budget paths.	Proceed when one repeatable buyer path emerges.
15-45	Source-grounded quality	Prototype retrieval/drafting on anonymized prior questionnaires.	Proceed if source coverage and draft acceptance are high enough to reduce expert load.
30-60	Live workflow pilot	Run two real packages with 3-5 design partners.	Proceed if cycle time and routing burden improve without unsafe release behavior.
45-75	Pricing	Ask pilot customers to convert at target annual pricing.	Proceed if willingness to pay supports human + model economics.
60-90	Onboarding repeatability	Onboard new account using a standardized process.	Proceed if setup does not depend on founder/manual research.
75-90	Expansion evidence	Observe demand for adjacent knowledge workflows.	Do not build unless pull is repeated across multiple target accounts.

ASSUMPTION LEDGER

The design is a portfolio of hypotheses, not a finished truth.

Assumption	Current confidence	What would falsify it
Target accounts face recurring questionnaire burden	Medium-High	Most qualified accounts handle fewer than ~10 significant packages annually.
RevOps/security coordination is painful enough to own budget	Medium	Problem exists but is treated as low-cost administrative work.
Approved historical answers are sufficient for grounded drafting	Medium	Novel questions dominate or source materials are too fragmented/stale.
Human review can be bounded	Medium	Every response requires extensive expert/legal interpretation.
\$18k median annual price is viable	Low-Medium	Buyers value labor savings but not enough to fund annual software.
Onboarding can become standardized	Medium	Each account requires extensive manual knowledge cleanup.
Model cost can remain below ~10% of revenue	Medium	High-context processing or premium models are required for most questions.
Expansion follows successful wedge adoption	Low	Customers want only transactional questionnaire automation.

Operating rule

Every material design decision should be tied to evidence and a trigger for re-evaluation. This prevents early customer requests, investor narratives, and technical possibility from silently becoming permanent business strategy.

RISK ARCHITECTURE

The largest risks are not all technical.

Risk	Failure mode	Design response
Accuracy / hallucination	Unverified answer leaves company and creates commercial/security risk.	Source grounding, confidence thresholds, named release owner.
Knowledge staleness	Previously approved answer becomes wrong after policy/product change.	Owner, validity date, refresh reminders, change history.
Security / confidentiality	Sensitive internal materials are mishandled.	Tenant isolation, access controls, scoped ingestion, audit trail.
Service creep	Team manually researches and authors too much content.	Track human minutes; price or reject non-product work.
Buyer ambiguity	No function clearly owns budget.	Qualification and buyer experiments before scaling sales.
Low frequency	Workflow is painful but too occasional for annual SaaS value.	Qualify by package volume and enterprise-sales intensity.
Platform drift	Founders expand to generic knowledge assistant before wedge repeatability.	Explicit expansion gates tied to retention and pull.

Risk principle

Trust controls are part of the business model because they influence buyer confidence, onboarding burden, support cost, sales cycle, and willingness to delegate work to the product.

EXPANSION LOGIC

A platform should emerge from proven adjacency, not from the pitch deck.

Expansion	Precondition	Evidence required
Security questionnaire knowledge	Initial wedge	Repeat usage + high source-grounded answer quality.
Trust center / evidence sharing	Customers repeatedly need the same external proof.	Repeated demand and clear owner/buyer overlap.
RFP response	Question structure and knowledge sources materially overlap.	Multiple accounts pull product into RFP workflow without major new service burden.
Legal / procurement responses	Approved language and risk controls can be bounded.	Legal review burden is predictable and source-groundable.
Internal knowledge assistant	Broader search need emerges from existing users.	Wedge retention strong; internal use has a distinct measurable job.
Adjacent industries	Workflow remains similar and sales/control burden does not change materially.	Pilot economics and implementation repeatability match core segment.

Expansion gate

No major adjacency should be approved because it increases TAM. It should be approved because customer behavior shows that the same knowledge asset, buyer relationship, and operating system can support another high-value workflow without breaking economics or trust.

SCENARIO ANALYSIS

Three businesses can be built from the same technology.

Scenario	Business design	Near-term benefit	Primary risk
A · Generic AI knowledge platform	Broad internal search + external response across departments.	Large conceptual market and flexible demos.	Weak urgency, crowded category, diffuse buyer, longer path to repeatability.
B · Questionnaire-response wedge	Narrow enterprise-sales workflow with source-grounded answers and approvals.	Clear trigger, measurable value, focused onboarding and GTM.	May appear smaller; requires discipline to resist adjacent requests.
C · Managed response service + software	Team provides substantial human completion/research alongside product.	Can create strong customer outcomes early and learn quickly.	Service economics may dominate and hide weak software repeatability.

Fellstead design view

Scenario B is the most coherent starting design for this illustrative case. It creates a narrow, observable transaction with a definable buyer, product boundary, pricing logic, and validation path. Scenario C can be used deliberately for learning only if human work is separately costed and bounded. Scenario A should remain an expansion thesis, not the launch model.

BUSINESS DESIGN BLUEPRINT

The proposed business on one page.

Layer	Design
Target customer	B2B SaaS / software companies, 50-500 employees, recurring enterprise questionnaires, limited specialist response capacity.
Trigger	Live enterprise deal creates security/procurement/compliance response package.
Buyer / user	Economic buyer: CRO/COO/security leader. Operational owner: RevOps, security, deal desk, or equivalent.
Core value	Faster verified response with less expert coordination and better answer consistency.
Product	Source-grounded AI response workspace + approved knowledge base + human approval workflow.
Human/AI boundary	AI retrieves/drafts; named humans approve sensitive/novel content and external release.
Pricing	Target \$18k-\$24k annual core plan + standardized onboarding.
GTM	Trigger-led founder sales + referral partners who see enterprise-response pain.
Economic constraint	AI + onboarding + support must preserve software-like gross contribution; manual knowledge service tracked separately.
Primary proof	Cycle-time reduction, draft acceptance, lower expert routing, annual conversion, repeatable onboarding.
Expansion rule	Add workflows only when repeated customer pull uses same knowledge asset and operating model.

FOUNDER / INVESTMENT SUMMARY

If only one page is used internally, use this one.

What is being built

Asterlane should begin as a controlled AI response workspace for enterprise security, procurement, and compliance questionnaires - not as a generic enterprise knowledge assistant.

Why this wedge

The workflow is recurring, externally triggered, tied to revenue, measurable, and dependent on verified company memory. Those properties create a clearer path to willingness to pay than a broad AI-assistant proposition.

What must be true

Target accounts experience enough questionnaire volume. A repeatable buyer owns the problem. Source-grounded drafting materially reduces expert coordination. Onboarding can be standardized. Buyers convert near target annual pricing. Model and support cost remain compatible with software economics.

Decision for the next 90 days

Do not build the platform vision. Build the smallest system necessary to test the workflow, buyer, pricing, trust model, and onboarding economics.

What would change the design

Low workflow frequency, weak annual willingness to pay, highly bespoke expert review, or onboarding that behaves like consulting would require a different model - potentially a service-assisted business, a different customer segment, or abandonment of the thesis.

NEXT VALIDATION QUESTIONS

Questions the business must answer before scale.

01. What percentage of target accounts process enough meaningful questionnaires each year to support annual SaaS pricing?
02. Which function actually feels the coordination cost strongly enough to become the operational owner?
03. Which buyer can approve \$18k-\$24k without a long enterprise procurement cycle?
04. What share of questions can be answered from existing approved knowledge with high-confidence sources?
05. How much human review remains after the knowledge base has matured?
06. What onboarding work is product setup versus customer-specific research or content creation?
07. How variable is model/API cost across short, long, and evidence-heavy questionnaires?
08. Does faster response measurably affect sales-stage duration or only internal convenience?
09. Which adjacent workflow is pulled by customers without requiring a different buyer or operating model?
10. What evidence would justify investing in a broad platform rather than deepening the initial wedge?

CASE NOTE AND LIMITATIONS

This is a fictional illustrative Business Design case. Asterlane AI, Inc. is not a real client. All names, figures, pricing, customer observations, economics, evidence, and recommendations are simulated. The design demonstrates how Fellstead can structure a business opportunity before significant capital, hiring, or software-development resources are committed. Actual Business Design work depends on the evidence available, the market, the company's constraints, and explicit client objectives. Fellstead provides business analysis and design, not legal, tax, accounting, audit, securities, or other regulated professional advice, and does not guarantee business outcomes.

fellstead.com · hello@fellstead.com