



AI products, services, and commercialization

Using AI to create differentiated offerings, customer experiences, and go-to-market capability.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 24 / EMBER & SKY

EXECUTIVE WORKING EDITION - 08 OCT 2026

AT A GLANCE

Using AI to create differentiated offerings, customer experiences, and go-to-market capability.

K

Most businesses first use AI to cut internal cost. The larger prize can be on the revenue side: new services, smarter product features, better customer education and sharper commercial packaging. It is also where mistakes are most visible, because customers see them. This playbook helps leaders decide which offerings to build, how to test demand before building, how to design quality, trust and pricing, how to measure commercial results, and how to launch in stages. It is practical guidance for service firms, product companies and hybrid businesses, with no promises about market outcomes.

WHO IT IS FOR

Founders, CEOs, product and service line leaders, marketing and sales leaders, and commercial strategy teams at small and mid-sized businesses considering AI-enabled offerings.

TIME TO USE

60 minutes to read; one working session to apply

∞

What you will have at the end

- A shortlist of AI-enabled offering ideas scored against customer value and feasibility
- A low-cost way to test demand and quality before heavy investment
- A pricing and packaging draft with unit economics you can check
- A staged launch plan with metrics and a stop rule

Contents

01	From efficiency to offerings customers will pay for	03
02	Four places AI can create commercial value	04
03	How AI-enabled offerings are built and where quality comes from	05
04	Who owns an AI-enabled offering	07
05	Trust, claims, data and contracts	08
06	Testing demand, quality and unit economics	10
07	A staged path from idea to launch	11
08	Common pitfalls and an illustrative scenario	13
09	Platform and partner choices	14
10	A first quarter that produces evidence	15
11	Terms used in this publication	16




EXECUTIVE BRIEF

From efficiency to offerings customers will pay for

01



Cost savings help the P&L once. A better offering can change how customers choose you.

Many businesses start with internal use of AI and stop there. That is sensible, but it leaves a second question unasked: what can customers now get from us that they could not before, or could not at this price, speed or level of personalization? Answering it can mean a new service line, a feature inside an existing product, a better way to teach customers to succeed with what they bought, or a sharper way to package and sell what you already do.

Commercial AI work differs from internal work in three ways. Customers experience the quality directly, so errors damage trust and brand. Customers compare, so a feature that every competitor can copy in a quarter is a weak moat. And unit economics matter: model usage, human review and support all cost money per customer, and a price set without them can lose money as usage grows. Start with the customer's problem, not the technology.

- Look for customer jobs that are slow, inconsistent, expensive or intimidating today.
- Prefer offerings that combine AI with your proprietary knowledge, data, relationships or process, which are harder to copy.
- Test demand and quality with real customers before committing to a build.
- Model unit economics early, including review, support and usage costs.
- Launch in stages, with a human fallback and clear limits on what the offering promises.

**HONESTY RULE**

Describe what the offering does and does not do in plain terms. Overclaiming about AI capabilities creates customer disappointment, support costs and, in some places, regulatory and consumer protection risk.



WORKING QUESTION

What do your customers still do by hand, wait for or avoid because it is too slow or too complicated?

WHERE IT FITS

Four places AI can create commercial value

02

Sort opportunities into four families so you compare like with like.

FAMILY	WHAT IT LOOKS LIKE	EXAMPLE QUESTION
New service design	A service made viable by faster or cheaper delivery, or a new tier of existing one	Which advisory deliverable could be offered as a fast, fixed-fee product?
Product features	Capabilities inside your product such as search, summaries, recommendations or drafting help	Which task do users abandon, and could a feature complete it for them?
Customer education and success	Guided onboarding, answers drawn from your documentation, personalized training	What questions flood support in the first 30 days?
Commercial packaging and go-to-market	Proposals, sales enablement, personalized content, pricing options and segment messaging	Where do deals slow because we cannot respond or tailor quickly enough?

Each family has a different risk and speed. Customer education and sales enablement are usually fastest and lowest risk, because a human can review output before it reaches the customer or because the assistant draws only from your approved content. Product features carry more engineering and quality risk. New service lines carry the largest strategic effect and the largest uncertainty about demand and delivery capacity.

Screening questions

- Customer value: would a named customer pay for this, or change supplier because of it?

- Differentiation: what do we have that others do not, such as data, expertise, trust, distribution or process?
- Feasibility: can we deliver with the data and skills we have, at consistent quality?
- Risk: what happens when it is wrong, and who bears the cost?
- Economics: does margin hold at the price customers will accept?


COMMON MISTAKE

Adding a chatbot because competitors did. A feature is worth building when it solves a specific customer problem better than the existing path.

WORKING QUESTION

Which one of your existing strengths would be hardest for a competitor to match with the same AI tools?


HOW IT WORKS

How AI-enabled offerings are built
and where quality comes from

03



Quality in a customer-facing offering comes from the system around the model: your content, your process and your review.

Most commercial AI offerings combine a general-purpose model with your own material. A customer asks a question; the system retrieves relevant passages from your knowledge base; the model drafts an answer from them; a rule or a person checks the output where the stakes require it. Differentiation lies in what you feed it, the workflow you wrap around it and the judgment you add, not in the model itself, which competitors can often access too.



01 Define the job

State the customer, the task, the output and what good looks like in measurable terms.

02 Assemble the proprietary layer

Gather the content, data, templates and expert judgment that make your answer better than a generic one.

03 Design the workflow

Decide what is automated, what is reviewed by a person, what is refused and how the customer reaches a human.

04 Test and tune

Run realistic scenarios, score quality, and fix content gaps and instructions before real customers see it.

Where quality problems come from

- Thin or outdated content, so the system has little of value to say
- Unclear scope, so the offering attempts tasks it cannot do well
- No refusal design, so it answers what it should decline
- Inconsistent results, so customers lose trust after one bad experience
- No feedback loop, so errors repeat

BUILD, BUY OR CONFIGURE

Many offerings can be configured from existing platforms rather than built from scratch. Choose the cheapest path that gives acceptable quality, and keep your proprietary content portable so you can change tools later.

WORKING QUESTION

For your best idea, what is the proprietary layer, and is it written down anywhere?

OPERATING MODEL

Who owns an AI-enabled offering

04

A commercial offering needs a product owner, a quality owner and a commercial owner, even in a small company.

ROLE	OWNS	TYPICAL PERSON
Offering owner	Customer problem, scope, roadmap, results	Product or service line leader
Quality owner	Test sets, review standards, error handling	Senior practitioner or operations lead
Commercial owner	Pricing, packaging, positioning, sales enablement	Marketing or sales leader
Content owner	The knowledge, templates and data the offering relies on	Subject matter expert
Risk and privacy lead	Customer data, terms, claims, regulatory fit	Legal, compliance or operations
Customer success	Onboarding, feedback, escalations	Support or account manager

In a small business one person may hold several roles. The point is that each responsibility has a name against it. Without a quality owner, nobody notices degradation. Without a content owner, the knowledge base ages. Without a commercial owner, the offering is built but never sold. Hold a short weekly review during launch, covering usage, quality scores, customer comments, costs and decisions needed.

Think carefully about delivery capacity. A new AI-enabled service line may attract demand faster than your experts can review output. Decide the review model in advance: whether every output is reviewed, a sample is reviewed, or only flagged cases. Price and staff accordingly. A fixed-fee offering that depends on unlimited expert review will erode margins as volume grows.

Operating basics

- Named offering, quality, commercial and content owners

- Documented scope: what is included, excluded and refused
- Review model with capacity to match expected volume
- Feedback channel for customers and a log of issues
- Weekly launch review in the calendar
- A defined route for customer complaints and refunds

WORKING QUESTION

If demand doubled in a month, which part of your delivery chain would break first?

RISKS AND CONTROLS

Trust, claims, data and contracts

05

Customers will forgive a limit that is honestly stated. They rarely forgive a failure they were told could not happen.

RISK	CONTROL	OWNER
Overstated claims about accuracy or capability	Review marketing copy against test results; state limits plainly	Commercial owner and risk lead
Wrong output harms a customer	Review model matched to stakes; clear human fallback; remedy process	Quality owner
Customer data misused or exposed	Collect minimum data; state use in terms; restrict access; check vendor terms	Risk and privacy lead
Intellectual property uncertainty over inputs or outputs	Check rights to the content you feed in; clarify ownership of outputs in contracts	Legal
Cost per customer exceeds price	Model unit economics; monitor usage; usage limits	Offering owner

RISK	CONTROL	OWNER
Unequal or unfair treatment	Test across customer segments; Quality owner review complaints	

Marketing statements about AI are scrutinized by consumers and regulators in many places. Keep claims specific and supportable: what the feature does, in which situations, with what limits. Avoid language implying certainty, human-level expertise or guaranteed results unless you can substantiate it. Keep the test evidence behind each claim on file.

Terms of service and contracts should address what data you collect, how you use it, whether it is used to improve the service, who owns outputs, what happens when the output is wrong, and the limits of your liability. Sector rules may add requirements, for example in financial, health or children's services. Have counsel review before launch. This playbook is not legal advice.

DISCLOSURE

Tell customers when they are interacting with an automated system or receiving AI-assisted work, and give them an easy way to reach a person. Check current disclosure and consumer protection requirements where you sell.

WORKING QUESTION

Which sentence in your current marketing could you not support with a test result?



MEASURING VALUE

Testing demand, quality and unit economics

06

Measure desirability, quality and economics separately. A strong result on one does not rescue a weak result on another.

Demand

Before building, talk to ten to twenty target customers. Describe the problem, show a mock-up or run a manual version of the service, and ask what they would pay and what they would stop doing to use it. Behavior beats opinion: a signed pilot agreement, a deposit or a booked trial is stronger evidence than enthusiasm. Record each conversation in the same format so you can compare.

Quality

Build a test set of 40 to 60 realistic customer scenarios with agreed good answers, including awkward and out-of-scope ones. Score each output as correct, partly correct, wrong or properly declined. Agree a pass mark before testing. Re-run after content or tool changes, and sample live outputs weekly.

METRIC	HOW TO MEASURE	NOTE
Qualified interest	Number of target customers who commit to a trial or pilot	Count commitments, not compliments
Activation and retention	Share who complete first use; share who return after 30 and 90 days	Compare against your non-AI baseline
Quality score	Test set results and live sampling	Track by scenario type
Customer satisfaction	Short survey after use; support themes	Include those who escalated to a human
Contribution margin per customer	Price minus model usage, review time, support and payment costs	Check at current and doubled volume
Sales cycle and win rate	Days to close and share won, for AI-supported proposals versus before	Many factors; look at the trend

ILLUSTRATIVE ARITHMETIC

Suppose a fixed-fee service costs 40 minutes of expert review per customer. If an expert hour costs your business a given amount, the review alone sets a price floor before usage and support. Use your own costs, and rerun the sum if volume grows.



WORKING QUESTION

What evidence of paying demand would you need to see before spending real money on a build?



IMPLEMENTATION

A staged path from idea to launch

07

Move from manual to assisted to automated only as the evidence earns it.



- 01 Frame**
Choose one offering idea. Write the customer, the problem, the promise and what you will not do.
- 02 Validate demand manually**
Deliver the service by hand for a few friendly customers, using AI behind the scenes with an expert reviewing. Learn what they value.
- 03 Design the system**
Capture the proprietary content, workflow, refusal rules and review model. Build a minimum version.





04 Beta

Release to a limited group under clear terms, with a human fallback, a feedback form and weekly reviews.

05 Price and package

Set tiers using margin data from the beta. Decide limits on usage, review and support included.

06 Launch and enable

Publish accurate claims, train sales and support, and prepare answers for common objections and security questions.



Launch readiness

- At least several customers have committed money or time
- Quality meets the agreed pass mark on the test set
- Unit economics hold at expected volume
- Claims, terms and disclosure are reviewed
- Support knows how to handle errors and refunds
- Stop and rollback criteria are written down



For commercial packaging inside the business, such as faster proposals, start with sales enablement. Build a library of approved case descriptions, service descriptions and pricing logic, and let an assistant assemble first drafts that a salesperson tailors. Track proposal turnaround time, the share of proposals using approved language and win rate. Keep prices and commitments with a human approver.



For customer education, begin with your existing documentation and the support tickets that recur most. An assistant that answers from approved content, shows its source and hands off to support when unsure can reduce repetitive questions and speed onboarding. Measure ticket volume by theme, time to first value and satisfaction.

WORKING QUESTION

Which customer could you serve by hand next month, with AI behind the scenes, to learn what they really value?

PITFALLS

Common pitfalls and an illustrative scenario

08



Offerings fail on demand, quality and economics, usually in that order of neglect.

- Building before talking to customers.
- Treating a general-purpose model as the differentiator.
- Pricing without counting review, support and usage costs.
- Launching with no refusal design or human fallback.
- Overclaiming in marketing and then handling the disappointment.
- Neglecting the content owner, so the knowledge ages.
- Chasing many offerings at once instead of proving one.



ILLUSTRATIVE SCENARIO (HYPOTHETICAL)

A small training provider designs an AI-assisted personalized learning plan as a paid add-on. It first delivers ten plans by hand with AI drafting and a trainer reviewing, and learns that customers value the trainer's check more than the speed. It prices the plan with the trainer review included and uses the drafts to cut preparation time. The lesson: the human layer was part of the product.



A related trap is competing on the feature rather than the outcome. If your competitors add the same capability next quarter, customers will compare on results, trust and service. Keep investing in the parts of the offering that are specific to you: your expertise, your data, your relationships and your standard of care.

Finally, resist treating early enthusiasm as proof. Early adopters tolerate rough edges. Test with a more typical customer before assuming the mainstream will behave the same way.



WORKING QUESTION

If a competitor copied your feature next quarter, what would customers still choose you for?

BUYING AND VENDORS

Platform and partner choices

09

Choose tools that let you ship quickly without trapping your content or your customers.

TOPIC	ASK	WHY IT MATTERS
Customer data	How is customer data stored, used, retained and deleted? Is it used to train models?	Trust and contract promises
Quality control	Can we test on our scenarios, see sources and set refusal rules?	Customer-facing quality
Cost model	How are usage costs charged, and what happens at ten times our volume?	Unit economics
Branding and control	Can we control wording, branding and escalation paths?	Customer experience
Portability	Can we export our content, configuration and logs?	Avoid lock-in
Reliability	What uptime, support and change notification do you commit to?	Customers notice outages

When using an implementation partner, define success in commercial terms: pilot customers signed, quality pass mark reached, cost per customer within target. Retain ownership of your content, your test sets and your customer relationships. Ask for documentation and training so your team can operate and improve the offering after launch.

Before you commit

- Pricing is modelled at current and growth volume
- Trial results on your own scenarios are documented
- Contract terms on data, support and change are reviewed
- You own your content, prompts and test sets
- Exit and export terms are agreed



WORKING QUESTION

What would it cost you, in time and money, to switch platforms after launch?



30 / 60 / 90

A first quarter that produces evidence

10

Use 90 days to choose one offering, test real demand and prove quality and economics before scaling.



First 30 days: ground it

- List customer jobs that are slow or painful and screen ideas against value, differentiation, feasibility, risk and economics
- Choose one offering and name the offering, quality, commercial and content owners
- Interview ten to twenty target customers and record commitments, not compliments
- Document the proprietary content and expertise the offering relies on



Days 31 to 60: test it

- Deliver the service manually with AI behind the scenes for a few friendly customers
- Build the minimum system with scope limits, refusal rules and a human fallback
- Create the test set and agree the pass mark
- Model unit economics, including review, usage and support costs

Days 61 to 90: decide

- Run a limited beta with weekly review of quality, cost and customer comments
- Set pricing, packaging and usage limits from beta data
- Review claims, terms and disclosure with counsel and prepare sales and support enablement
- Decide to launch, extend the beta, change course or stop, and write down why

GLOSSARY

Terms used in this publication

11

TERM	MEANING
Unit economics	Revenue and cost for each customer or unit sold, showing whether growth adds profit.
Contribution margin	Price minus the variable costs of serving a customer.
Proprietary layer	The content, data, process and expertise that make your offering better than a generic one.
Differentiation	The reasons customers choose you over alternatives.
Beta	A limited release to a small group to learn before wide launch.
Activation	A new customer completing the first meaningful use of the offering.
Retention	The share of customers who keep using or paying over time.
Packaging	How an offering is bundled, tiered and priced.
Sales enablement	The content and tools that help sales staff respond and close deals.
Refusal design	Deliberate behavior that declines out-of-scope or unsafe requests.

TERM	MEANING
Human fallback	A clear route for a customer to reach a person.
Test set	A fixed group of realistic scenarios with agreed good answers used to score quality.

K'

A note on use

This playbook is general business guidance. It does not guarantee commercial results and is not legal or financial advice. Check consumer protection, privacy and sector rules where you sell.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

II

