FLASHFEED STUDY GUIDE

AI for Small Business

Study guide by FlashFeed

November 2026 edition · pack version 0.1.18

flashfeedlearningapp.com

AI for Small Business: study guide

A beginner's practical guide to choosing useful AI workflows, prompting with context, researching carefully, supporting customers, automating safely, protecting data, and running a measured 30-day pilot.

Edition	November 2026 (2026-11-21)
Pack version	0.1.18
Contents	8 lessons · 244 cards in the app · 97 self-test questions · 8 pictures
In the app	flashfeedlearningapp.com/packs/ai-for-small-business

How to use this guide

Take one lesson at a time: read the study guide, then go over any key facts. Where a lesson has a self-test, try it from memory before you check the answer key at the back. Follow the Go deeper links when you want more.

Free, with no ads and no account to make.

Contents

01	AI Foundations for an Owner	4
02	Prompting as Business Communication	10
03	Research, Data, and Decisions	16
04	Marketing, Sales, and Customer Service	22
05	Operations, Automation, and Agents	28
06	High-Stakes Work: Money, Contracts, and People	34
07	Privacy, Security, and Responsible Use	40
08	A 30-Day Implementation Playbook	46
	Answer key	52
	Sources	60
	Image credits	61
	Useful links	62
	About FlashFeed	63

LESSON 1 OF 8

AI Foundations for an Owner

What AI Can and Cannot Do · Start with a Workflow, Not a Tool

WHAT YOU'LL BE ABLE TO DO

- Distinguish productive AI assistance from fluent but unsupported output, and verify material claims before acting.
- Choose a frequent, low-risk workflow for a first AI pilot instead of beginning with a tool or high-stakes decision.

Capabilities

Generative AI is best understood as a highly capable **pattern machine**. It can draft, summarize, classify, compare, brainstorm, translate, and reorganize information at unusual speed. That makes it useful, but it does not make it an oracle. A polished answer can still be incomplete, stale, or wrong.

For a small-business owner, the most useful distinction is between **generation** and **analysis**. Generation creates a new draft: a customer reply, job description, checklist, promotion, or meeting agenda. Analysis looks for patterns or implications in material you provide: a set of reviews, a sales export, a contract, or a collection of support tickets. Both require review. The model may misunderstand a column, invent a fact, or smooth over uncertainty — confidently producing invented, unsupported, or incorrect information despite fluent-sounding wording. This failure mode is called a **hallucination**, and it is why every material claim needs verification no matter how polished the answer sounds.

AI becomes more useful when you supply the local truth: your actual offer, customer, policy, price, examples, and constraints. No general model knows your business as well as the people operating it. Treat your expertise as the source and AI as a fast collaborator.

The adoption statistics are a useful lesson in evidence. The U.S. Census Bureau reported overall business AI use around 17–20% from late 2025 through spring 2026 under its survey wording. A U.S. Chamber small-business survey reported 58% generative-AI use in 2025. Those figures are not interchangeable: they involve different populations, definitions, and methods. Before repeating a headline number, ask who was surveyed and what counted as use.

Workflow First

The best place to begin is not a product comparison. It is a **workflow**: a repeatable sequence of inputs, decisions, actions, and outputs. List recurring work and look for a task that is frequent, structured, easy to review, and low consequence when the first draft is wrong.

Good early candidates include drafting routine replies, turning notes into a checklist, creating variations from approved marketing copy, summarizing internal meetings, or organizing customer feedback. Poor

first candidates include final hiring decisions, unsupervised payments, legal commitments, or anything where one error can materially harm a person or the business.

Start by **augmenting** the existing process. Run the AI-assisted version beside the current method and compare time, quality, and correction effort. This parallel period teaches you the exceptions that a generic demonstration never shows. Only after the workflow is understood should you consider additional automation.

The owner remains responsible for defining good work. If a process is inconsistent or undocumented, AI will reproduce the inconsistency faster. Clean up the process, define what acceptable output looks like, and then test assistance.

A practical owner exercise

Choose three tasks from yesterday. For each, write down the input, the judgment required, the output, and the consequence of an error. A routine follow-up email probably has a clear input and a reversible output. Deciding whether to fire an employee has ambiguous inputs, legal consequences, and a decision that should never be delegated to a general AI tool. The exercise trains you to see task characteristics rather than product features.

Then test one safe task twice. In the first attempt, provide only a short request. In the second, include the business context, a real example, and a definition of good output. Compare the results. This usually demonstrates an important principle: model capability matters, but the surrounding information and process often matter just as much.

Warning signs in AI advice

Be cautious when an answer presents precise savings, adoption rates, or revenue gains without a traceable source and definition. Be cautious when a product recommendation ignores your existing software, staff skills, data sensitivity, or workflow. Be especially cautious when anyone promises effortless passive income, fully autonomous operations, or guaranteed outcomes.

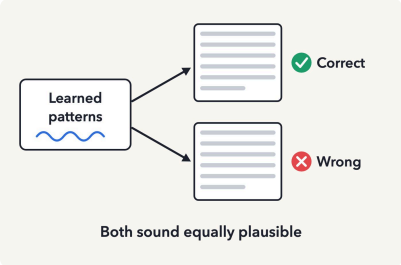
Useful AI advice makes tradeoffs visible. It identifies what the owner must supply, what needs checking, what could go wrong, and how to test value with limited exposure. That standard applies to this pack too: product names are examples, while the durable lesson is how to evaluate work.

Key terms

Hallucination
A confident-sounding AI output that contains invented, unsupported, or incorrect information.

Workflow
A repeatable sequence of inputs, decisions, actions, and outputs used to complete a business task.

Key facts



FlashFeed original diagram · CC-BY-4.0

A Pattern Machine, Not an Oracle

A language model generates plausible responses from learned patterns. It can be remarkably capable, but fluent wording is not proof that a claim is true.

Useful Does Not Mean Deterministic

The same request can produce different answers. Use each AI output as a draft or recommendation whose reliability depends on the task, context, and review process.

You Supply the Local Truth

AI usually does not know your current prices, policies, customers, inventory, or judgment. Your business context is what turns a generic answer into useful work.

Workflow First

Start with a recurring business problem, then evaluate tools. Buying an AI subscription before naming the job often creates cost without adoption.

The Best First Task

A strong first task is frequent, structured, easy to check, and low consequence if the first draft is wrong.

Augment Before You Automate

Use AI alongside the existing process before allowing it to act alone. A parallel run exposes errors, exceptions, and missing instructions.

Matched pairs

Each item on the left goes with the one on the right.

Generation → Creating a new draft from instructions and context

Analysis → Finding patterns or implications in information you provide

High-fit first workflow → Frequent, structured, reviewable, and low risk

Low-fit first workflow → Rare, ambiguous, irreversible, or high stakes

Step by step

Run a safe first AI experiment

1. Choose a low-risk task you already understand
2. Define what a good result looks like
3. Give the AI context and a clear desired output
4. Compare the result with work you trust
5. Record what required correction

Early experiments should teach you where AI helps and where review is still essential.

Build an AI opportunity list

1. List tasks repeated at least weekly
2. Estimate time spent on each task
3. Mark tasks that follow a recognizable pattern
4. Rate the cost of an error
5. Choose one frequent, low-risk task to test

A task inventory keeps adoption grounded in work rather than novelty.

Self-test

Answer each question from memory, then check your work against the answer key.

1. A polished AI answer includes three statistics but no sources. What should you do first?

- A Publish it because the prose is confident
- B Replace the numbers with larger estimates
- C Verify each statistic against reliable sources
- D Ask the AI whether it is certain

2. In AI terms, what is a 'hallucination'?

- A A slowdown caused by too many simultaneous requests
- B A confident-sounding output that contains invented or unsupported information
- C A setting that shortens the length of a response
- D A deliberate refusal to answer a sensitive question

3. Which task is an example of AI analysis rather than generation?

- A Finding patterns in a batch of customer reviews you upload
- B Writing a customer reply from your notes
- C Brainstorming names for a new promotion
- D Drafting a job description from a short brief

4. Which statement about a polished AI answer is NOT correct?

A. A polished answer can still be incomplete or stale
B. Material claims should be verified before acting
C. The same request can produce different answers
D. Fluent, confident wording confirms the claim is accurate

5. Why can't the Census Bureau's roughly 17 to 20 percent figure and the Chamber's 58 percent figure be treated as interchangeable?

A. They measured entirely different decades
B. They used different populations, definitions, and methods
C. One counted only failed AI projects
D. Both figures were later retracted as errors

6. Two surveys report very different AI-adoption rates for small businesses. What should you check before citing either number?

A. The population surveyed and how 'use' was defined
B. Which of the two figures more news articles repeat
C. Whether the outlet is well known
D. Whether a chatbot confirms the same figure when asked

7. An AI-generated report states a precise revenue increase with no traceable source. What should you do before acting on it?

A. Trust it because the number is precise
B. Ask a colleague if it sounds plausible
C. Verify it against an independent source first
D. Round the number down slightly

8. Which is the strongest first AI workflow for a neighborhood service business?

A. Auto-send AI replies to every customer complaint
B. Draft a once-a-year strategic plan
C. Compare AI tools before listing any tasks
D. Draft weekly social posts from approved offers

9. How is a 'workflow' defined?

A. A ranked list of AI vendors to compare
B. A one-time decision about which software to buy
C. The total monthly cost of an AI subscription
D. A repeatable sequence of inputs, decisions, actions, and outputs

10. Which task is the WORST candidate for a first AI pilot?

A. Turning meeting notes into a checklist
B. Deciding whether to fire an employee
C. Summarizing an internal meeting
D. Drafting routine customer replies

11. What is the main benefit of running AI alongside your existing process before automating it?

- A It exposes errors, exceptions, and missing instructions
- B It lets you skip writing instructions
- C It speeds up the work immediately, before any testing
- D It removes the need to define good output

12. A strong first AI task should be:

- A Complex and requiring expert judgment
- B Rare, ambiguous, and legally binding
- C Frequent, easy to check, and low consequence if wrong
- D Irreversible but high in value

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/ai-foundations-for-an-owner

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

LESSON 2 OF 8

Prompting as Business Communication

Give AI Context and a Clear Target · Iterate, Critique, and Reuse

WHAT YOU'LL BE ABLE TO DO

- Write an AI brief with a clear task, relevant context, the audience, constraints, and the expected output.
- Improve weak output through critique, added examples, and reusable prompt-library assets rather than repeated one-shot prompting.

Context

A prompt is not a magic phrase. It is a short business brief. Useful requests usually contain five ingredients: the **task**, the relevant **context**, the audience, constraints, and a description of the expected output.

Compare “write a customer email” with a real brief: “Draft a 120-word reply to a customer whose order is delayed. Use the shipping policy below. Acknowledge the inconvenience, do not promise a delivery date, offer tracking and cancellation as next steps, and use our calm, direct tone.” The second request is better because the work is defined.

Examples are often stronger than adjectives. An approved email can show rhythm, warmth, terminology, and level of detail better than “sound professional.” Before sharing examples, remove personal, confidential, or unnecessary information.

Define **acceptance criteria** before reading the answer. Criteria might include required facts, maximum length, a specified format, prohibited claims, links that must be checked, or a policy that must be followed. Criteria make review faster and reduce the temptation to approve an answer merely because it sounds polished.

For complex work, let the AI ask questions before drafting. A proposal, plan, or customer communication may depend on details you forgot to include. “Ask up to five questions before you begin” turns missing context into a visible step.

Iteration

Treat the first response as a draft. Compare it with the criteria, identify the most important weakness, provide missing context or an example, and request a focused revision. A working conversation usually outperforms a single elaborate command.

Separate the **draft pass** from the **critique pass**. First create a workable version. Then ask for weaknesses against explicit criteria, unsupported claims, missing edge cases, or language that does not match the audience. Finally, perform the human edit. AI should help you see options; it should not make you less attentive.

When a workflow succeeds, save more than the final output. A useful **prompt library** stores the instructions, source material, good examples, acceptance checklist, and owner. Date the entry because tools, policies, and business facts change.

Credible practical educators reinforce this conversational approach. Ethan Mollick's current guides emphasize supplying context, using AI on work you understand, and checking errors. OpenAI Academy's small-business materials use a similar Goal–Context–Output–Boundary checklist. The durable idea is not a secret prompt; it is clear communication plus informed review.

Build a request from real work

Start with a task you can judge, such as replying to a new inquiry. Gather one approved reply, the current service area, the response-time promise, and the questions that must be answered before quoting. Ask the AI to draft a response using only that information. Require it to mark missing details rather than invent them.

Review the draft in layers. First check factual accuracy and policy. Next check whether the response advances the customer's next step. Then edit tone and wording. Reviewing substance before style prevents a common failure: polishing a message whose underlying promise is wrong.

If the output disappoints, diagnose the gap. A wrong fact may mean the source material was missing or ignored. Generic language may mean the audience and examples were vague. A bad structure may mean the expected format was undefined. An unsafe promise may mean the constraints were incomplete. Diagnosis produces a reusable improvement; asking for “better” does not.

Maintain the library

Each saved workflow should have a name, owner, purpose, approved inputs, prohibited data, current instructions, examples, checklist, and review date. Test it when a model changes or a business policy changes. Archive instructions that no longer work.

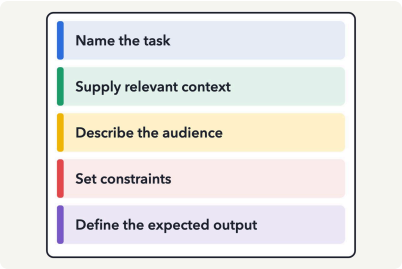
Do not treat a prompt library as a bag of spells copied from social media. Its value comes from encoding your business knowledge and review standard. A prompt that works for a dentist, restaurant, or HVAC company may fail for your customer, risk, and voice.

Key terms

Acceptance criteria
Specific conditions an output must satisfy before it is considered usable.

Prompt library
A maintained collection of tested instructions, examples, and checklists for recurring work.

Key facts



FlashFeed original diagram · CC-BY-4.0

Prompt Like a Brief

A useful request names the task, supplies relevant context, describes the audience, sets constraints, and defines the expected output.

Examples Beat Adjectives

A real email, proposal, or response often communicates your voice more clearly than instructions such as 'make it professional.' Remove sensitive details before sharing examples.

Define Done

Acceptance criteria make quality review faster. Specify required facts, length, format, tone, prohibited claims, and what the output must help someone do.

Use a Working Conversation

Strong results usually emerge through questions, critique, and revision. Treat the first answer as material to improve, not a final deliverable.

Let the AI Ask Questions

When important context is missing, instruct the AI to ask a few questions before drafting. This is especially useful for proposals, plans, and customer communications.

Save the Pattern, Not Just the Output

When a workflow succeeds, save the instructions, source documents, examples, and review checklist. That package is more reusable than one finished draft.

Matched pairs

Each item on the left goes with the one on the right.

Task → The job you want the AI to perform

Context → Business information needed to perform the task well

Draft pass → Create a workable first version

Critique pass → Identify weaknesses against explicit criteria

Step by step

Write a reusable AI brief

1. State the task
2. Describe the audience and business context
3. Attach or paste trusted source material
4. Set constraints and approval boundaries
5. Define the output format and acceptance criteria

Store successful briefs so the team does not reinvent them every time.

Improve an AI draft

1. Compare the draft with acceptance criteria
2. Point out the most important gap
3. Provide missing context or an example
4. Request a focused revision
5. Perform the final human edit and approval

Change one important variable at a time so you learn what improves the result.

Self-test

Answer each question from memory, then check your work against the answer key.

1. Which request is most likely to produce a usable customer email?

- A Write something friendly about our policy
- B Draft a 120-word reply to a delayed-order complaint using the policy below; acknowledge the delay, avoid promising a date, and offer two next steps
- C Act as the world's best marketer
- D Write an email

2. A useful AI request usually contains which five ingredients?

- A Task, context, audience, constraints, and expected output
- B Model, temperature, token limit, seed, and format
- C Greeting, body, closing, subject line, and attachment
- D Keywords, tone, length, deadline, and signature

3. Specific conditions an output must satisfy before it is considered usable are called:

- A The draft pass
- B A prompt library
- C Acceptance criteria
- D Context

4. Why include an approved example email in a brief instead of only saying 'sound professional'?

A. It lets the AI skip reading the constraints
B. It guarantees the output needs no human review
C. It uses fewer tokens than writing adjectives
D. It shows rhythm, warmth, terminology, and level of detail better than adjectives

5. Acceptance criteria might specify all of the following EXCEPT:

A. A policy that must be followed
B. Required facts the output must include
C. A maximum length or specified format
D. The model's temperature and random seed

6. An AI reply has the right facts but arrives as one long paragraph when you needed a table. What was most likely undefined in the brief?

A. The constraints on promises
B. The expected output format
C. The business context
D. The audience for the reply

7. An AI draft is generic and misses your differentiator. What is the best next move?

A. Publish it and improve next month
B. Switch tools immediately
C. Add specific business context and a good example, then revise
D. Ask for more enthusiasm

8. A maintained collection of tested instructions, examples, and checklists for recurring work is called a:

A. Prompt library
B. Critique pass
C. Acceptance criterion
D. Draft pass

9. During the critique pass of an AI conversation, what should you do?

A. Create the first workable version
B. Identify weaknesses against explicit criteria
C. Delete the draft and start over
D. Publish the output to save time

10. When a workflow succeeds, what is most valuable to save?

A. The instructions, source material, examples, checklist, and owner
B. The AI's temperature setting
C. A screenshot of the chat window
D. Only the final polished draft

11. A draft makes an unsafe promise to a customer. Using the diagnosis approach, what was most likely incomplete?

A. The maximum length
B. The expected output format
C. The constraints
D. The approved example emails

12. A prompt finally produces reliable customer emails. What should you save so the workflow is reusable?

- A Nothing; regenerate from scratch each time
- B Only the final email you liked
- C A screenshot of the AI conversation
- D Instructions, sources, examples, and checklist together

13. Which practice keeps a saved prompt-library entry useful as models and policies change?

- A Sharing prompts without an owner listed
- B Copying popular prompts from social media
- C Dating entries and retesting when policy or model changes
- D Deleting the checklist after the first success

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/prompting-as-business-communication

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

LESSON 3 OF 8

Research, Data, and Decisions

Research Current Markets and Competitors · Analyze Data and Make Better Decisions

WHAT YOU'LL BE ABLE TO DO

- Verify current market claims with primary sources and triangulation before using them in a business decision.
- Separate observations from inferences and check data quality before recommending an action.

Research

AI can compress the mechanical parts of research: generating questions, locating sources, comparing claims, summarizing long documents, and organizing a brief. It cannot remove the need to evaluate evidence.

Use current search when the answer could have changed: prices, regulations, product features, market conditions, competitors, schedules, and public officials. Ask for links and open the important ones. For consequential claims, prefer **primary sources** such as government data, an original survey, a contract, vendor documentation, or direct customer evidence.

Distinguish a source from a citation-shaped decoration. A link must actually support the claim. Check the publication date, population, geography, definition, and method. If a statistic would change your decision, verify it independently.

Use **triangulation** for uncertain questions. Customer interviews, transaction data, search behavior, competitor observations, and public statistics each reveal different pieces. Agreement among independent methods is more useful than several articles repeating the same original claim.

A decision-ready research brief states the decision, separates evidence from assumptions, records source dates, identifies uncertainty, and explains what new evidence would change the recommendation. Honest uncertainty is a feature.

Data Decisions

Before analysis, inspect the data. Missing values, duplicate records, inconsistent date formats, mixed units, and unclear column definitions can produce elegant nonsense. Ask the AI to describe what the file contains and what appears ambiguous before requesting conclusions.

Separate **observations** from **inferences**. “Average order value fell in June” may be supported directly. “The new promotion attracted worse customers” is an explanation that requires more evidence.

For forecasting, use scenarios. A base, upside, and downside case reveal the assumptions carrying the result. **Sensitivity analysis** changes one important input at a time to show whether the recommendation is robust or fragile.

Finish with a compact decision memo: the decision and deadline, verified observations, assumptions, realistic options, recommendation, risks, and the next checkpoint. Keep the source data available so a reviewer can challenge the interpretation.

Example: researching a new service

Suppose a landscaping company is considering a subscription maintenance plan. Begin with the decision: whether to pilot the plan, for which customers, and at what minimum margin. Ask AI to help build a research checklist rather than estimate the answer from general knowledge.

Evidence might include current customer requests, cancellation reasons, job frequency, travel time, labor capacity, competitor offers, and local seasonality. Public competitor pages can show positioning, but not whether those plans are profitable. Customer interviews reveal interest, but stated interest is weaker than a deposit or signed pilot. Internal job records may be strongest for estimating cost.

Ask the AI to create a table that labels every entry as observed data, sourced external evidence, assumption, or calculation. Require a source and date for external claims. Then ask what missing fact has the greatest chance of reversing the recommendation.

Review numbers deliberately

Check totals independently. Confirm units, date ranges, tax treatment, and whether averages hide important customer segments. A monthly average may conceal a seasonal cash shortage. Revenue may look attractive while travel and rework destroy margin.

When comparing scenarios, change explicit assumptions such as customer count, labor hours, retention, or material cost. Do not let the AI quietly change several assumptions at once. The goal is to understand the decision, not merely produce a sophisticated spreadsheet.

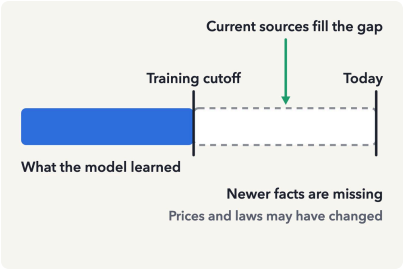
Finally, make a small test when possible. A limited offer to a defined customer group often provides better evidence than another week of broad online research.

Key terms

Triangulation
Checking a material claim against multiple independent sources or methods before relying on it.

Sensitivity analysis
Testing how a result changes when an important assumption or input changes.

Key facts



FlashFeed original diagram · CC-BY-4.0

Current Questions Need Current Sources

For prices, laws, competitors, products, and market conditions, use an AI mode that searches the web and provides citations. Model memory alone may be stale.

Prefer Primary Sources

For material decisions, trace important claims to government data, original research, contracts, vendor documentation, or direct customer evidence.

Definitions Change the Number

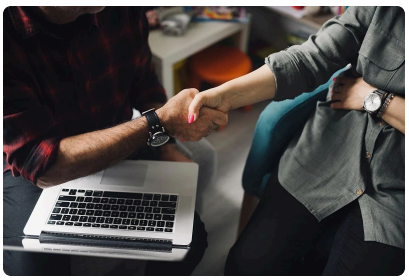
AI adoption estimates can differ because surveys ask different populations and define 'using AI' differently. The U.S. Census measured about 17–20% of businesses using AI from December 2025 through May 2026, while other surveys using looser definitions of 'using AI' report figures several times higher.

Inspect Before You Analyze

Before asking for conclusions, identify missing values, duplicate rows, inconsistent dates, unclear units, and columns whose meaning is uncertain.

Separate Observation from Inference

'Returns increased in June' is an observation. 'The new supplier caused the increase' is an inference that needs additional evidence.



Wikimedia Commons · Shixart1985 · CC BY 2.0

Use Scenarios, Not One Forecast

A base case, upside case, and downside case reveal which assumptions drive the decision and where a contingency plan is needed.

Matched pairs

Each item on the left goes with the one on the right.

Primary source → Original data, law, policy, documentation, or firsthand evidence

Secondary source → A summary or interpretation of someone else's evidence

Observation → A pattern directly supported by the available data

Inference → An interpretation or explanation drawn from observations

Step by step

Create a decision-ready research brief

1. Write the decision the research must support
2. List the claims that would change the decision
3. Gather current sources and record their dates
4. Separate evidence, assumptions, and estimates
5. Verify the highest-impact claims independently
6. Summarize implications and remaining uncertainty

A research brief should make uncertainty visible rather than hiding it behind polished prose.

Build a one-page decision memo

1. State the decision and deadline
2. Summarize verified observations
3. List assumptions and uncertainties
4. Compare two or three realistic options
5. Recommend an action with reasons
6. Define the next measurement or checkpoint

Keep the evidence visible so a reviewer can challenge the recommendation.

Self-test

Answer each question from memory, then check your work against the answer key.

1. An AI says local demand rose 40% this year. What evidence would be strongest?

- (A) A second AI repeating the claim
- (B) Current local sales, search, survey, or government data with a defined comparison period
- (C) A vendor blog with no methodology
- (D) A confident explanation of why demand rose

2. Which is the strongest example of a primary source for a material business decision?

- (A) A news article quoting an industry analyst
- (B) A competitor's marketing page describing a trend
- (C) An original customer survey your business conducted
- (D) A blog post summarizing a government report

3. Which action best demonstrates triangulation when validating a demand claim?

A. Averaging two guesses produced by the same model
B. Reading three articles that all cite the same original study
C. Asking the AI to restate its estimate more confidently
D. Comparing customer interviews, transaction data, and public statistics

4. Which question is LEAST likely to require an AI mode that searches the web for current sources?

A. The meaning of a standard arithmetic operation
B. The current interest rate on a business loan
C. A new product's latest features
D. A competitor's current pricing

5. The U.S. Census found about 17 to 20% of businesses using AI, while another survey using a looser definition reported a much higher figure. Why do these differ so much?

A. One of the two figures is simply a typo
B. The surveys were conducted in different countries
C. Adoption collapsed sharply between the two surveys
D. The surveys asked different populations and defined using AI differently

6. You need to confirm a local demand claim before relying on it. Which approach best reflects triangulation?

A. Reading the same blog post twice
B. Cross-checking sales data, interviews, and a government survey
C. Asking the AI to restate the number more confidently
D. Averaging three AI-generated guesses

7. Why is agreement among independent research methods more convincing than several articles repeating one claim?

A. More articles always mean more accuracy
B. It takes less time to find
C. Repeated articles often share one unverified source
D. Independent methods can each fail in the same way

8. What should happen before AI recommends actions from an uploaded sales spreadsheet?

A. Inspect data quality and clarify column meanings
B. Remove the lowest month
C. Convert every value to a percentage
D. Ask for the most surprising conclusion

9. Which statement is an inference rather than an observation?

A. Total revenue was lower than in May
B. Average order value fell in June
C. The new supplier caused the increase in returns
D. Returns increased in June

10. What does a sensitivity analysis do?

- A Removes outliers before any analysis begins
- B Tests how a result changes when one important input changes
- C Produces a single most-likely forecast number
- D Averages several months together to smooth the trend

11. What is the main benefit of using base, upside, and downside cases instead of one forecast?

- A It reveals which assumptions drive the decision
- B It guarantees a more accurate prediction
- C It removes the need to check data quality
- D It hides uncertainty from a reviewer

12. Before asking AI for conclusions, which is NOT part of inspecting data quality?

- A Checking for duplicate records
- B Clarifying ambiguous column definitions
- C Identifying missing values
- D Choosing the chart color scheme

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/research-data-and-decisions

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

LESSON 4 OF 8

Marketing, Sales, and Customer Service

Create Marketing Without Losing Your Voice · Support Sales and Customer Service

WHAT YOU'LL BE ABLE TO DO

- Use AI to adapt marketing content while preserving brand voice and substantiating claims.
- Design customer-facing assistance that uses approved knowledge and escalates unknown, sensitive, or exceptional cases.

Marketing

AI is excellent at variation. It can turn one verified source into channel-specific drafts, but it needs a **brand brief**: audience, promise, tone, proof points, preferred terms, prohibited claims, and examples of good work.

Build from a **content atom**—an original customer story, guide, demonstration, interview, or case study. Extract the core lesson and evidence, then adapt it into an email, social posts, a sales handout, or video scripts. Adaptation matters: a LinkedIn post and a customer email serve different contexts.

Keep claims truthful. AI may supply plausible performance promises, testimonials, or comparisons that the business cannot support. The owner must verify facts, permissions, links, and disclosures before publication.

Human contribution matters for intellectual property as well as quality. The U.S. Copyright Office explains that AI-assisted work can be protected when people contribute sufficient expressive authorship, while prompts alone generally do not establish authorship over a purely generated output. Use AI as part of a human creative process and preserve drafts and edits when authorship matters.

Sales Service

Customer-facing AI should be grounded in approved information: current hours, products, pricing rules, service areas, policies, and canonical answers. General model knowledge is not a substitute for your business records.

AI can help score inbound leads, but **lead qualification**—checking whether a prospect matches your defined needs and buying conditions—still needs a person to weigh judgment calls the data cannot capture.

Design the handoff before launch. An **escalation rule** tells the workflow when to stop and transfer the case. Common triggers include uncertainty, threats, legal questions, sensitive personal information, unusual exceptions, or refunds above a defined limit.

Start in draft-only mode. Have the AI propose replies while a person reviews every one. Test normal questions, incomplete questions, emotional complaints, attempts to override policy, and requests the system cannot fulfill.

Personalization should be relevant and expected. Use information the customer knowingly supplied for the relationship. Do not reveal hidden profiling or unrelated data merely because a system can access it.

Review conversation logs for recurring questions, wrong answers, failed handoffs, and policy gaps. The purpose is not to eliminate people; it is to give routine questions fast answers while preserving human attention for exceptions and relationships.

A simple brand brief

A practical brief can fit on one page. Name the primary customer, the problem you solve, the promise you can support, three proof points, the desired tone, preferred terminology, phrases to avoid, and two approved examples. Include boundaries such as “never imply same-day service” or “do not discuss a customer's account in public review responses.”

Use the brief to critique content, not merely generate it. Ask whether a draft makes an unsupported claim, sounds interchangeable with a competitor, ignores the reader's situation, or lacks a useful next step. AI is often more valuable as a tireless editor than as a source of endless new posts.

Test customer workflows safely

Create a set of test cases before enabling customer-facing automation. Include an easy question, missing information, an angry customer, a request outside policy, a privacy-sensitive question, and a prompt attempting to override the assistant's rules. Record the correct answer or handoff for each case.

Review whether the system admits uncertainty. A chatbot that says “I don't have enough information, so I'll connect you with the owner” can preserve trust. A chatbot that invents a warranty term can create a dispute.

Measure more than response speed. Track resolution, escalation accuracy, corrections, customer satisfaction, and whether staff must repair conversations later. Fast wrong answers are not an efficiency gain.

When using customer stories, reviews, photos, or messages in generated content, confirm permission and avoid exposing information that was provided for a different purpose.

Key terms

Content atom
A substantial source asset that can be adapted into multiple smaller channel-specific pieces.

Escalation rule
A defined condition that stops automation and transfers a case to human review.

Key facts

Give the Brand a Memory

A compact brand brief should define audience, promise, tone, proof points, prohibited claims, preferred terms, and examples of good work.

Create Once, Adapt Carefully

One useful customer story or guide can become an email, several social posts, a sales handout, and short video scripts when each version is adapted to its channel.

Human Authorship Still Matters

The U.S. Copyright Office says AI-assisted work can qualify for protection when humans contribute sufficient expressive authorship, but prompts alone generally do not make a purely AI-generated output copyrightable.

Ground Customer Answers

A customer assistant should answer from approved policies, product information, hours, pricing rules, and canonical answers—not from general model knowledge.

Design the Handoff

Automation is safer when it knows when to stop. Uncertainty, threats, legal questions, sensitive personal information, unusual exceptions, and refunds above a defined limit should route to a person.

Personalize with Restraint

Useful personalization reflects relevant customer context. Creepy personalization reveals data the customer did not expect you to use.

Matched pairs

Each item on the left goes with the one on the right.

Brand voice → The consistent personality and language choices used in communication

Proof point → Specific evidence supporting a marketing claim

Lead qualification → Checking whether a prospect matches defined needs and buying conditions

Escalation → Routing a case to a person with the authority or expertise to resolve it

Step by step

Turn one source into a small campaign

1. Choose an original source with a clear customer lesson
2. Extract the main message and supporting proof
3. Draft channel-specific versions
4. Check facts, claims, links, and brand voice
5. Have an owner approve external content
6. Measure response and record what worked

Repurposing should preserve the core truth while adapting format and emphasis.

Prepare a customer-answer workflow

1. Collect the most common questions
2. Write and approve canonical answers
3. Identify information the assistant may access
4. Define prohibited promises and escalation triggers
5. Test normal, ambiguous, and adversarial questions
6. Review logs and update the knowledge regularly

Start in draft-only mode before allowing automatic replies.

Self-test

Answer each question from memory, then check your work against the answer key.

1. AI drafts an advertisement claiming your service is 'guaranteed to double sales.' What should you do?

- (A) Remove or substantiate the claim and review it for legal and brand risk
- (B) Replace 'guaranteed' with 'almost certain'
- (C) Publish it if it sounds persuasive
- (D) Ask AI to add a testimonial

2. Which of the following is NOT typically part of a one-page brand brief?

- (A) Preferred terminology and phrases to avoid
- (B) Three proof points that support your promise
- (C) A month-by-month sales revenue forecast
- (D) The desired tone of voice

3. You wrote one detailed customer success story. Following the content-atom approach, what should you do with it?

A. Post it only on the one channel where it performs best
B. Adapt it into an email, social posts, a sales handout, and video scripts, each tailored to its channel
C. Discard it after a single use to keep your content fresh
D. Publish the identical text on every channel to stay consistent

4. According to the U.S. Copyright Office, when can AI-assisted work qualify for copyright protection?

A. Whenever the prompt is detailed enough to guide the output
B. Never, because AI was involved at some stage
C. Automatically, because all AI output is protected
D. When a human contributes sufficient expressive authorship

5. In a brand brief, what is a 'proof point'?

A. A channel-specific version of a content atom
B. A phrase the brand has agreed never to use
C. Specific evidence that supports a marketing claim
D. The overall personality and language of the brand

6. A chatbot cannot find an approved answer to a customer's warranty question. What should it do?

A. End the conversation without explanation
B. Invent the most likely policy
C. Give a generic industry answer
D. Say it is unsure and escalate to a person

7. What is an escalation rule in a customer-service workflow?

A. A defined condition that stops automation and transfers a case to human review
B. A cap on how many messages a customer may send per day
C. A script that lets the chatbot answer harder questions on its own
D. A ranking of prospects by how likely they are to buy

8. Which of these is LEAST likely to be set as an escalation trigger?

A. A customer making a legal threat
B. A routine question about current store hours
C. A refund request above your defined limit
D. A request involving sensitive personal information

9. A customer-facing AI assistant should base its answers primarily on what?

A. The general knowledge built into the language model
B. The most popular answers found on competitor websites
C. Whatever reply sounds the most confident and helpful
D. The business's approved policies, pricing, hours, and canonical answers

10. Which use of personalization in a customer reply is appropriate rather than 'creepy'?

- A Bringing up data the customer never expected you to have
- B Revealing a hidden profile the system quietly built about them
- C Referencing the order the customer just placed with you
- D Mentioning an unrelated purchase the customer made elsewhere

11. A customer assistant is being set up to answer product and policy questions. What should its answers be grounded in?

- A The AI's general training knowledge
- B Approved policies, product details, hours, and pricing rules
- C Competitor websites
- D Whatever sounds most helpful

12. Which source should a customer-facing AI treat as authoritative when answering a pricing question?

- A The business's own approved pricing rules and records
- B The customer's own estimate
- C A confident guess based on similar businesses
- D A general model's memorized knowledge

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/marketing-sales-and-customer-service

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

LESSON 5 OF 8

Operations, Automation, and Agents

Document Operations and Repetitive Work · Automate Carefully and Introduce Agents

WHAT YOU'LL BE ABLE TO DO

- Capture an experienced worker's real process, including happy paths and exceptions, as a testable standard operating procedure.
- Distinguish rule-based automation from an AI agent and place human approval before consequential actions.

Operations

Much of a small business lives in the owner's head. AI can help turn demonstrations, voice notes, meeting transcripts, and rough checklists into usable operating documentation.

Capture the real process. Narrate a task while doing it or interview the person who performs it. Ask AI to organize the material into purpose, inputs, steps, decisions, outputs, safety controls, and exceptions. The experienced worker—not the model—is the authority.

A **standard operating procedure** includes the happy path and the exception paths. What happens when a field is blank, a customer objects, an item is unavailable, or the primary system fails? Clean lists that omit exceptions create brittle automation.

Test the procedure with a different person. Observe where they pause, guess, or need undocumented access. Revise it, assign an owner, and set a review date. Documentation is a maintained operating asset.

AI meeting notes are useful when they distinguish decisions, action items, owners, deadlines, and unresolved questions. Always tell participants when recording or transcription is used and follow applicable policy and law.

Automation

Rule-based automation follows a predefined trigger and action. An **AI agent** can choose steps and use tools toward a goal. Greater flexibility creates greater need for boundaries.

Start with permissions. Decide what a workflow may read, create, change, send, spend, or delete. Use least privilege: provide only the access needed. Put consequential actions—payments, external messages, record changes, account changes—behind an **approval gate**.

Keep logs showing inputs, outputs, actions, errors, and approvals. Include a kill switch so a responsible person can stop the workflow immediately. Pilot with limited data and a limited audience before expanding.

Move in stages: document, simplify, draft, compare, control, then automate. This sequence feels slower than a flashy demonstration and is much faster than repairing an automation that quietly performed the wrong task at scale.

Find the hidden decisions

When documenting a process, ask the experienced worker where they pause, inspect, call someone, or make an exception. Those moments often contain the real expertise. “Send the estimate” may hide decisions about minimum margin, travel distance, available capacity, customer history, and required approval.

Ask AI to turn a transcript into a draft procedure, then mark every uncertain rule with a question instead of filling the gap. Resolve those questions with the process owner. Add screenshots or examples only when you have the right to share and maintain them.

Choose the right level of automation

Not every improvement needs an agent. A template may be enough. A rule-based integration may reliably copy form submissions into a CRM. AI becomes useful when the workflow must interpret variable text, classify an exception, summarize information, or draft a response.

Use a permissions table with rows for read, create, edit, send, spend, and delete. For each, state whether the workflow is prohibited, allowed, or allowed only after approval. Begin with read and draft permissions. Expand only when measured performance justifies it.

Test failure conditions: missing files, duplicate events, conflicting records, unavailable services, unusually large amounts, and malicious instructions hidden in incoming content. Confirm that retries do not send duplicate messages or payments.

Assign an owner who checks logs and exceptions. Automation without ownership becomes invisible operational debt.

Key terms

Standard operating procedure

A maintained description of how to perform a recurring task, including steps, responsibilities, controls, and exceptions.

Approval gate

A required human confirmation before a workflow performs a consequential action.

Key facts

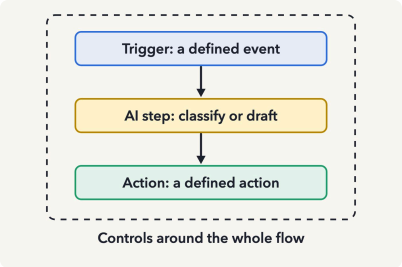
> **Capture the Real Process**
>
> The fastest way to draft a procedure is often to narrate the task while doing it, then have AI organize the transcript into steps, decisions, and exceptions.

Exceptions Are Part of the Procedure

A clean happy-path checklist is incomplete. Good operating documentation explains what to do when information is missing, a customer objects, or a system fails.

Meetings Need Owners and Deadlines

AI summaries become operationally useful when decisions, action items, owners, due dates, and unresolved questions are separated clearly.



FlashFeed original diagram · CC-BY-4.0

Automation Starts with a Trigger

A basic automation follows a rule: when a defined event occurs, perform a defined action. AI can classify or draft inside that flow, but the surrounding controls still matter.

Agents Need Boundaries

An agent may choose steps and use tools toward a goal. Limit what it can read, change, send, spend, or delete, and require approval for consequential actions.

Keep an Audit Trail

Logs make it possible to see what ran, what information was used, what the AI proposed, what action occurred, and who approved it.

Matched pairs

Each item on the left goes with the one on the right.

Happy path → The normal sequence when nothing unusual happens

Exception path → Instructions for unusual conditions or failures

Rule-based automation → A predefined trigger and action sequence

AI agent → A system that can select steps and use tools toward a goal within permissions

Step by step

Create and test an AI-assisted procedure

1. Observe or record the task being performed
2. Draft the steps, decisions, inputs, and outputs
3. Add safety checks and exception paths
4. Have an experienced person review it
5. Ask a different person to follow it
6. Revise and assign an owner and review date

A procedure is not finished until someone else can use it successfully.

Move a workflow toward automation

1. Document the current workflow
2. Remove unnecessary steps and clarify rules
3. Run AI in draft-only mode
4. Measure errors and identify exceptions
5. Add permissions, approval gates, and logs
6. Increase autonomy gradually and keep a kill switch

Do not automate a process you cannot explain or evaluate.

Self-test

Answer each question from memory, then check your work against the answer key.

1. What is the strongest source for a new procedure?

- A The actual process demonstrated by an experienced worker
- B A generic AI answer about the industry
- C An old checklist no one follows
- D A competitor's marketing page

2. In a standard operating procedure, what does the exception path describe?

- A The normal sequence followed when nothing unusual happens
- B Instructions for unusual conditions or failures
- C A required human confirmation before spending money
- D A predefined trigger and matching automatic action

3. Which of these is NOT something a well-formed operating procedure should include?

A. Exception paths for missing information or failures
B. An assigned owner and a review date
C. Steps, decisions, and safety controls
D. A guarantee that no experienced worker input is ever needed

4. When AI drafts a procedure from a transcript, what should you do with rules you are unsure about?

A. Let the model fill the gap with its best guess so the draft stays complete
B. Ship the draft and fix problems only after users report them
C. Mark each uncertain rule with a question and resolve it with the process owner
D. Drop the uncertain step to keep the happy path clean

5. After drafting a procedure, how should you check that it actually works?

A. Have the author reread it for missing steps
B. Test it with a different person and watch where they pause or guess
C. Ask the AI to check its own draft for gaps
D. Time how long the author takes to follow it

6. Which action should normally require approval in a beginner business automation?

A. Extracting dates from a document
B. Drafting an internal summary
C. Sending a refund from the business account
D. Labeling an email for review

7. What distinguishes an AI agent from rule-based automation?

A. It can choose steps and use tools toward a goal
B. It follows one predefined trigger and one fixed action
C. It only copies form submissions into a CRM
D. It never requires any permissions or boundaries

8. What is an approval gate in an automated workflow?

A. A rule that triggers an action automatically with no review
B. The least amount of access a workflow is granted
C. A log recording inputs, outputs, actions, and errors
D. A required human confirmation before a consequential action

9. Following least privilege, which permissions should a beginner automation usually start with?

A. Unlimited access to every connected system
B. Delete and account-change permissions
C. Full send and spend permissions
D. Read and draft permissions

10. For which task is simple rule-based automation usually enough, with no AI agent needed?

A. Summarizing a long meeting transcript
B. Classifying an ambiguous customer exception
C. Copying new form submissions into a CRM
D. Drafting a tailored reply to unusual feedback

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/operations-automation-and-agents

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

LESSON 6 OF 8

High-Stakes Work: Money, Contracts, and People

Use AI Around Money, Contracts, and Taxes · Support Hiring and People Work Fairly

WHAT YOU'LL BE ABLE TO DO

- Use AI for preparation around contracts while reserving material decisions for a qualified human.
- Use AI for hiring administration without delegating employment judgment or ignoring adverse-impact risk.

Finance Legal

AI can be valuable around high-stakes work without becoming the final authority. It can summarize a lease, compare versions, extract invoice fields, categorize transactions, draft questions, or explain unfamiliar language. The final filing, commitment, payment, or professional judgment stays with a qualified person.

For documents, ask for cited sections and direct quotations kept short enough to inspect. Compare the summary with the source. A **material term** can affect cost, rights, obligations, risk, or the decision to sign; unusual or unclear material terms belong on a question list for counsel.

For financial work, reconcile to original records. An attractive dashboard is not proof that invoices, statements, payroll, and ledger entries agree. Use AI to surface anomalies and organize review, then let the owner, bookkeeper, CPA, or responsible employee resolve them.

Do not rely on general AI for complex tax answers or calculations. Current law, jurisdiction, business structure, elections, and records matter. Use AI to prepare records and questions, not to manufacture certainty.

People

Existing employment law applies when software assists recruiting, assessment, scheduling, pay, promotion, or discipline. A vendor's AI label does not transfer the employer's responsibility.

Begin with documented, job-related criteria. AI may help draft a posting, remove jargon, organize interview notes, or map applications to objective requirements. It should not make final employment decisions or infer sensitive traits.

The EEOC and Department of Justice warn that algorithmic tools may screen out people with disabilities. Provide a clear accommodation process and an alternative when a tool creates a barrier.

Meaningful human review is more than clicking approve. The reviewer needs authority, time, relevant information, and the ability to challenge the recommendation. Document the criteria and final

decision. For location-specific requirements, consult current federal, state, and local guidance or qualified counsel.

Prepare, do not pretend

For a contract, ask AI to produce a table with the clause, plain-language summary, obligation, deadline, financial exposure, termination rule, and question for counsel. Require section references. This creates a review aid; it does not determine whether the contract is fair or enforceable.

For bookkeeping, AI can propose categories and flag unusual transactions. A responsible person should review supporting records and reconcile accounts. For cash-flow planning, identify which inputs are historical facts and which are assumptions. Never allow a forecast to authorize a transfer or borrowing decision automatically.

For tax work, use official forms, instructions, and a qualified preparer where appropriate. Product demonstrations and general chat answers may omit elections, exceptions, jurisdictional rules, or recent changes.

Keep hiring criteria visible

Write the role's essential duties and job-related requirements before reviewing people. Use the same documented criteria consistently. If AI summarizes applications, preserve access to the original material and allow reviewers to correct omissions.

Do not ask AI to infer personality, health, family status, age, race, religion, or other protected or sensitive characteristics. Avoid proxy criteria that are unrelated to performance. Provide candidates with a way to request accommodation or an alternative process.

Periodically compare who advances and who is screened out. A neutral-looking process can still produce harmful or unlawful outcomes. If a tool cannot be explained, tested, or challenged enough for the decision, do not use it there.

Key terms

Material term

A contract provision important enough that it could affect cost, rights, obligations, risk, or the decision to sign.

Adverse impact

A substantially different negative outcome for a protected group caused by a seemingly neutral employment practice.

Key facts

First Pass, Not Final Authority

AI can summarize a contract, organize receipts, draft questions, or flag unusual terms. It should not be treated as the final authority on legal, accounting, lending, or tax decisions.

Reconcile Against the Source

Financial summaries should be tied back to statements, invoices, payroll records, or the accounting system. A plausible total is not a reconciled total.

Use Professionals Where Judgment Matters

An attorney, CPA, bookkeeper, insurance professional, or lender can evaluate facts and obligations that a general AI assistant may miss.

Existing Employment Law Still Applies

Using an algorithm does not remove an employer's obligations. AI-assisted recruiting or assessment can still create unlawful discrimination.

Plan for Accommodation

The EEOC advises employers to consider whether algorithmic tools screen out people with disabilities and to provide a process for reasonable accommodation.

Keep Meaningful Human Review

AI can organize applications against documented job requirements, but a responsible person should review context, challenge errors, and make employment decisions.

Matched pairs

Each item on the left goes with the one on the right.

AI first pass → Summarize, classify, compare, or prepare questions

Qualified human decision → Accept responsibility for advice, filing, approval, or commitment

Objective criterion → A job-related requirement defined before reviewing candidates

Subjective judgment → An interpretation that can be influenced by assumptions or bias

Step by step

Use AI for high-stakes preparation

1. Use an approved tool and remove unnecessary sensitive data
2. Ask for a structured summary with cited sections
3. List assumptions, missing information, and questions
4. Compare important figures or clauses with original records
5. Send the source and questions to the responsible professional
6. Record the final human decision

Never let a summary replace the underlying document or record.

Use AI as a hiring aid

1. Define job-related requirements and evaluation criteria
2. Remove criteria unrelated to successful performance
3. Use AI for drafting or organization, not final decisions
4. Check outputs for unsupported assumptions and disparate treatment
5. Provide accommodation and alternative-process information
6. Document the human review and decision

Consult current federal, state, and local requirements for your location and workforce.

Self-test

Answer each question from memory, then check your work against the answer key.

1. AI flags an unusual indemnity clause in a commercial lease. What is the best next step?

- A Ask AI to guarantee the lease is safe
- B Ignore it because AI sometimes makes mistakes
- C Delete the clause yourself and sign
- D Use the flag to prepare questions for a qualified attorney

2. In a contract, what best describes a 'material term'?

- A A minor formatting detail that has no bearing on the agreement
- B A provision important enough to affect cost, rights, obligations, risk, or the decision to sign
- C Any clause written in legal boilerplate language
- D The physical materials a signed paper contract is printed on

3. Which statement about using AI for financial work is INCORRECT?

- A A bookkeeper or CPA should resolve flagged discrepancies
- B Summaries should be tied back to statements, invoices, payroll, or the ledger
- C A plausible-looking dashboard total proves the accounts actually agree
- D AI can help surface anomalies and organize a review

4. Why should a small business avoid relying on general AI for complex tax answers?

- A Current law, jurisdiction, structure, elections, and records all matter
- B Tax questions are too simple to need any preparation
- C Tax software is always cheaper than any AI tool
- D AI cannot read any tax forms or instructions

5. Which task is appropriate for AI's first pass rather than a qualified human's final decision?

- A Summarizing a lease and drafting questions for counsel
- B Approving a loan commitment on the owner's behalf
- C Accepting responsibility for a tax filing
- D Authorizing a payment based on a forecast

6. What is the safest beginner use of AI in recruiting?

- A Make the final hiring decision
- B Draft a job description from documented role requirements for human review
- C Automatically reject applicants based on personality predictions
- D Rank candidates by names and photos

7. What is 'adverse impact' in employment?

- A A drop in overall company hiring during a slow year
- B A deliberate, openly stated policy of excluding a group
- C Any negative feedback an employee receives during a review
- D A substantially different negative outcome for a protected group from a seemingly neutral practice

8. The EEOC and Department of Justice warn that algorithmic hiring tools may do what?

- A Raise the cost of each new hire
- B Expose applicant data to outside vendors
- C Screen out people with disabilities
- D Slow down the hiring process

9. What does 'meaningful human review' of an AI hiring recommendation require?

- A Authority, time, relevant information, and ability to challenge the recommendation
- B Trusting the tool's output because it looks objective
- C Reviewing only the candidates the tool already advanced
- D Clicking approve quickly to keep the process moving

10. Which task should you NOT ask an AI hiring tool to do?

- A Remove jargon from a job posting
- B Organize interview notes for a human reviewer
- C Map applications to documented, job-related requirements
- D Infer a candidate's health, age, or family status

11. Why should an owner periodically compare who advances versus who is screened out by a hiring tool?

- A To reduce recruiting costs
- B To speed up the review process
- C To detect a disparate negative outcome for a protected group
- D To satisfy the vendor's warranty terms

12. The EEOC has warned that algorithmic hiring tools carry which specific risk?

- A Slower processing times
- B Screening out qualified people with disabilities
- C Reduced applicant volume
- D Higher advertising costs

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/high-stakes-work-money-contracts-and-people

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

LESSON 7 OF 8

Privacy, Security, and Responsible Use

Protect Customer and Business Data · Build Human Oversight and Quality Controls

WHAT YOU'LL BE ABLE TO DO

- Classify business information, minimize shared data, and check vendor retention and training controls before use.
- Match oversight to consequence using risk tiers and human review.

Data

Before sharing information with an AI service, name the purpose and classify the data. Public marketing copy is different from a customer list, payroll file, health record, trade secret, or contract under confidentiality.

Practice **data minimization**. Share only fields needed for the task. Remove direct identifiers and unnecessary narrative. De-identification is not magic; combinations such as location, dates, and transaction details can still identify someone.

Evaluate the exact product and account type. Training, retention, sharing, administrator access, integrations, and deletion controls can differ between personal and business offerings. For example, as of 2026, OpenAI states that its Business, Enterprise, and API offerings do not use customer inputs and outputs for model training by default. That fact applies to those offerings, not every AI service, and vendor policies can change, so reverify current terms before relying on this.

The FTC advises businesses to collect only what they need, protect it, honor privacy promises, and dispose of it securely. If your privacy policy says one thing and your AI workflow does another, the new technology does not excuse the mismatch.

Governance

Use consequences to set controls. Internal brainstorming may need ordinary review. Customer messages need factual and brand review. Payments, employment decisions, legal commitments, safety advice, and sensitive-data workflows need stronger testing, approval, and documentation.

NIST's voluntary AI Risk Management Framework uses four functions: **Govern, Map, Measure, and Manage**. For a small business, that can be simple. Assign an owner and policy; understand the use case and affected people; test performance and risks; then choose controls and monitor results.

A **human in the loop** must be placed where review can still change the outcome. Add approval gates before consequential actions, escalation rules for uncertainty, least-privilege access, and audit logs.

Plan for incidents. Know how to stop the workflow, preserve records, correct bad messages or data, notify the owner, and prevent recurrence. Reassess when the vendor, model, source data, or business

process changes.

Create a simple data map

List the information used by the proposed workflow, where it comes from, where it is sent, who can access it, how long it is kept, and how it is deleted. Include connected apps: an AI service may be acceptable while a plug-in or automation sends data somewhere unexpected.

Classify data using a small set of labels such as public, internal, confidential, and restricted. Define examples for your business. Customer contact details, employee records, payment information, health information, credentials, and trade secrets deserve explicit treatment.

Ask vendors concrete questions. Is content used for training by default? How long is it retained? Can administrators access it? Which subprocessors receive it? What happens when an account is closed? Can connected tools take actions? Product names and plan names change, so record the date and link for each answer.

Scale controls with risk

Create three practical tiers. Low-risk work is internal, reversible, and uses nonsensitive data. Medium-risk work reaches customers or influences operations but remains reviewable. High-risk work affects money, rights, safety, legal obligations, employment, or sensitive data.

For each tier, define required review, testing, approval, logging, and incident response. Prohibit uses whose risk you cannot manage. A short policy that people understand is stronger than a long document no one follows.

Train staff to report surprising output, accidental data sharing, incorrect customer messages, and unauthorized actions without hiding mistakes. Fast reporting limits harm and improves the system.

Key terms

Data retention
How long a service or business keeps information before deletion or archival.

Human in the loop
A workflow design in which a person reviews, approves, corrects, or can stop AI-supported work at a meaningful point.

Key facts



AI-generated (gpt-image-1.5)

Share the Minimum

Do not paste an entire customer file when the task only needs a product category and date. Data minimization reduces exposure and often improves focus.

Plan Labels Are Not Data Controls

Free, personal, business, and enterprise products can have different training, retention, sharing, and administration settings. Verify the current terms and controls for the exact product you use.

Business Products May Offer Stronger Defaults

For example, OpenAI states (as of 2026) that ChatGPT Business, Enterprise, and API data are not used to train models by default. Every vendor's terms should be verified separately and periodically before sensitive use.



AI-generated (gpt-image-1.5)

Control by Consequence

The more an AI output can affect money, rights, safety, reputation, or sensitive data, the stronger the testing, approval, documentation, and monitoring should be.

Govern, Map, Measure, Manage

NIST organizes AI risk work into four functions: establish governance, understand the context, assess risks and performance, and prioritize responses.

Plan for Failure

An AI incident plan should say how to stop the workflow, preserve logs, correct affected records or messages, notify responsible people, and prevent recurrence.

Matched pairs

Each item on the left goes with the one on the right.

Public data → Information intentionally available to anyone

Confidential data → Restricted business, customer, employee, or partner information

Reversible action → An action that can be corrected with limited harm or cost

Consequential action → An action that can materially affect money, rights, safety, or trust

Step by step

Check data before using an AI tool

1. Name the business purpose
2. Classify the information involved
3. Remove fields that are not necessary
4. Verify the tool's current training, retention, access, and sharing controls
5. Confirm customer, employee, contractual, and legal permissions
6. Choose a safer method or stop if controls are insufficient

When in doubt, use synthetic or de-identified examples for experimentation.

Run an AI workflow preflight

1. Define the intended use and owner
2. Identify affected people, data, and decisions
3. Rate the consequence of errors or misuse
4. Set tests, approval gates, permissions, and logs
5. Pilot with limited scope
6. Monitor outcomes and document incidents and changes

Reassess the workflow when the model, vendor, data, or business process changes.

Self-test

Answer each question from memory, then check your work against the answer key.

1. You want AI to segment a customer list. What should happen first?

(A) Ask the AI provider whether breaches are impossible

(B) Upload the full list to a free tool

(C) Replace customer names, but keep every other identifier

(D) Verify authorization, tool controls, needed fields, and retention before sharing data

2. What does practicing data minimization mean when using an AI service?

- A Keep every record indefinitely in case it is useful later
- B Share all available data so the AI has full context
- C Share only the fields actually needed for the task
- D Encrypt data before pasting it into any tool

3. Your team upgrades from a free to a business AI plan. What should you assume about its training and retention settings?

- A They are identical across all plan tiers
- B They may differ, so verify the current terms for that exact product
- C Business plans never retain any data, regardless of settings
- D Free plans always offer stronger privacy

4. Using the labels public, internal, confidential, and restricted, which item is public data?

- A A marketing brochure published on your website
- B An employee payroll file
- C A signed contract under confidentiality
- D A customer contact list

5. In vendor and data terms, what does 'data retention' refer to?

- A The number of people who can access a dataset
- B The process of encrypting stored records
- C Whether inputs are used to train a model
- D How long a service or business keeps information before deletion or archival

6. What does data minimization mean when sharing information with an AI tool?

- A Convert all data into images before sharing
- B Share only the fields needed for the task
- C Delete the AI vendor's account after use
- D Encrypt every file before upload

7. You only need a product category and date for an AI task. What should you do before sharing customer data?

- A Share names but keep every other identifier
- B Ask the vendor to retain the data forever
- C Remove fields the task does not need
- D Upload the entire customer file for context

8. Before using an AI vendor for sensitive work, what should you check?

- A Default training use and data retention period
- B How many five-star reviews the tool has
- C Whether the tool offers a free trial period
- D Whether the vendor offers phone support

9. Which OpenAI product tiers are stated to skip customer inputs and outputs for model training by default?

- A Free, Plus, and Pro personal accounts
- B Business, Enterprise, and API
- C Plus, Pro, and Enterprise
- D Every OpenAI plan, including Free

10. Which workflow needs the strongest review and controls?

- A Generating internal brainstorming ideas
- B Reformatting a meeting agenda
- C Drafting three subject-line options
- D Automatically changing employee pay based on an AI score

11. In AI governance, which best describes a consequential action?

- A Any step a person happens to review before it runs
- B An action that can be corrected later, with limited harm or cost
- C An action that can materially affect money, rights, safety, or trust
- D A task that stays entirely inside internal brainstorming

12. Where should a human in the loop be placed to be effective?

- A At a point where review can still change the outcome
- B After the action is complete and irreversible
- C Only during the initial software setup
- D Wherever it slows the workflow the least

13. Which action is NOT part of a sound AI incident plan?

- A Deleting all logs to reduce liability
- B Stopping the workflow
- C Notifying the responsible people
- D Correcting affected records or messages

14. A cafe uses AI to brainstorm in-house promotion ideas that never leave the shop and are easy to redo. Which risk tier fits best?

- A High risk
- B Medium risk
- C Prohibited use
- D Low risk

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/privacy-security-and-responsible-use

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

LESSON 8 OF 8

A 30-Day Implementation Playbook

The 30-Day Pilot · Your Small-Business AI Operating System

WHAT YOU'LL BE ABLE TO DO

- Plan a measured 30-day pilot with a baseline, acceptance criteria, parallel run, and evidence-based expansion decision.
- Recognize the policies, approval gates, and audit logs that make AI use dependable.

Pilot

Choose one workflow for a 30-day **pilot**. A pilot is a limited implementation with a defined owner, scope, period, controls, and success criteria. It exists to create evidence, not to prove that AI is exciting.

In the first week, build a **task inventory** and choose a frequent, structured, reviewable, low-risk workflow. Record a **baseline**: current time, cost, error rate, turnaround, backlog, or customer outcome. Define acceptance criteria and the data the pilot may use.

In the second week, create the brief, examples, and checklist. Run a **parallel test** beside the current process. Record where the AI succeeds, where it fails, what context is missing, and how long review takes.

In the third week, use the workflow with real but limited work. Keep human approval. Track acceptance rate, correction time, errors, and exceptions. Do not expand scope in the middle merely because one result is impressive.

In the fourth week, compare results with the baseline. Include subscription cost, setup time, review time, time saved, avoided cost, and outcome changes in the **return on investment** calculation. Decide to stop, adjust, or expand.

Operating System

Turn a successful pilot into an owned capability. Write a short **AI use policy** covering approved tools, prohibited data, review rules, external use, incidents, and responsibility.

Maintain a **prompt library** with tested instructions, examples, source materials, criteria, owner, and review date. Add an **approval gate** before consequential actions, an **escalation rule** for uncertainty and exceptions, an **audit log** for accountability, and a **kill switch** for rapid shutdown.

Review the operating system periodically. Check access, vendor terms, source documents, failures, value, and whether the workflow still solves the right problem. Retire workflows that no longer earn their complexity.

For continued learning, prioritize educators who show full workflows, discuss limitations, and distinguish evidence from sponsorship. Ethan Mollick's One Useful Thing offers practical, research-aware guidance. NIST, the SBA, SBDCs, SCORE, and regulator resources provide durable guardrails. OpenAI Academy and other major vendors provide useful product-specific demonstrations; treat those as training for a tool, not neutral rankings of the market.

A weekly scorecard

Use a one-page scorecard throughout the pilot. Record volume processed, total time, AI cost, review time, accepted outputs, corrected outputs, serious errors, exceptions, and an outcome that matters to the business. Notes should explain unusual weeks.

Calculate value conservatively. Time saved is valuable only when it frees capacity, reduces delay, improves quality, or replaces a real cost. Ten minutes saved on work no one needed is not a return. Include setup and maintenance time, not just the subscription.

At the end, choose one of three decisions. Stop when the workflow does not create enough value or the risk is unacceptable. Adjust when the idea is sound but instructions, data, controls, or scope need work. Expand when results repeatedly meet criteria and the owner can support the operating burden.

Choosing people to learn from

Thought leaders can add practical examples, but use a filter. Prefer people who show inputs and outputs, identify failure cases, update old advice, cite research, and disclose sponsorships. Be skeptical of channels built around daily tool rankings, guaranteed income, or dramatic statistics without methodology.

For practical general guidance, Ethan Mollick is a strong starting point. For business education, use SBA, SCORE, and Small Business Development Center programs. For risk, use NIST, FTC, EEOC, and relevant state regulators. For specific products, use the vendor's current documentation and academy resources.

The goal is not to become an AI news expert. It is to build one useful, controlled workflow and develop the judgment to repeat the process.

Key concepts

Task inventory

- A list used before selecting an AI tool.
- It focuses attention on recurring work and bottlenecks.
- It records frequency, time, structure, and error consequence for candidate workflows.

Baseline

- A reference point established before a change.
- It prevents teams from relying only on impressions of improvement.
- It records current time, cost, errors, or outcomes for comparison with a pilot.

Pilot

- A limited implementation designed to produce evidence.
- It reduces exposure while a workflow is tested and refined.
- It uses a defined scope, owner, period, and success criteria before wider rollout.

Parallel run

- A transition technique that preserves an existing process temporarily.
- It makes side-by-side quality and timing comparisons possible.
- The new AI-assisted workflow operates alongside the manual workflow before replacement.

Acceptance criteria

- A quality tool defined before reviewing results.
- It makes approval less dependent on mood or impressive wording.
- It lists the conditions an output must satisfy to be considered usable.

Return on investment

- A comparison between value received and resources spent.
- For a pilot, its inputs can include subscriptions, setup, review time, savings, and avoided cost.
- It helps decide whether to stop, adjust, or expand an AI workflow.

AI use policy

- A governance document sized to the business.
- It gives staff common boundaries before a problem occurs.
- It names approved tools, prohibited data, review rules, external-use expectations, and incident reporting.

Prompt library

- A reusable knowledge asset rather than a single conversation.
- It helps a team repeat successful AI-assisted work.
- It stores tested instructions, examples, source materials, criteria, an owner, and a review date.

Approval gate

- A control placed before consequential action.
- It limits autonomy without preventing drafting or analysis.
- It requires a person to confirm an action before money moves, a message sends, or a record changes.

Audit log

- A record that supports troubleshooting and accountability.
- It helps reconstruct what a workflow did and why review failed.
- It records inputs, outputs, actions, timestamps, errors, and approvals.

Escalation rule

- A boundary for uncertainty, exceptions, or authority.
- It prevents a system from bluffing its way through unusual cases.
- It defines when automation stops and a case transfers to a responsible person.

Kill switch

- An operational safeguard for a malfunctioning workflow.
- It should be quick to use and known to responsible staff.
- It immediately stops automated actions while an incident is investigated.

Key facts

> **One Pilot Beats Ten Trials**
> A focused 30-day pilot creates evidence. Tool hopping creates opinions without a stable process or baseline.

> **Measure Before and After**
> Record current time, cost, error rate, turnaround, and customer impact before changing the workflow. Without a baseline, improvement becomes guesswork.

> **Track Usable Output**
> A useful pilot metric is the share of AI outputs accepted after review, along with the time and type of corrections required.



AI-generated (gpt-image-1.5)

> **Keep the First Policy Short**
> A beginner AI policy can fit on one page: approved tools, prohibited data, review requirements, external-use rules, incident reporting, and an owner.

> **Every Workflow Needs an Owner**
> Someone must maintain instructions, permissions, source material, tests, and performance—even when the workflow runs automatically.

> **Review the System, Not Just the Output**
> Periodically review tool access, vendor terms, failed outputs, time savings, source documents, prompts, and whether the workflow still solves the right problem.

Self-test

Answer each question from memory, then check your work against the answer key.

1. When should a small business expand an AI pilot?

- A After buying an annual plan
- B As soon as the output looks impressive
- C When competitors mention the same tool
- D When measured results meet criteria and risks are controlled

2. What best describes a 30-day pilot?

- A A permanent rollout across every team at once
- B A demo meant to show that AI is exciting
- C A time-boxed implementation run under clear ownership and success criteria
- D An open-ended trial with no fixed end date

3. Which measurement is NOT part of recording a baseline before the pilot?

- A The final expansion decision
- B Current time and cost
- C Error rate and turnaround
- D Backlog or customer outcome

4. In the second week, how should the AI workflow be tested against the current process?

- A By expanding scope before any review
- B By running a parallel test beside the current process
- C By replacing the current process entirely
- D By skipping the current process to save time

5. When building the first-week task inventory, which workflow should you choose for the pilot?

- A Whichever task competitors are automating
- B A rare, one-off task with high stakes
- C A frequent, structured, reviewable, low-risk workflow
- D The most creative and unpredictable task available

6. In a 30-day AI pilot, when should acceptance criteria be defined?

- A After results are already compared to competitors
- B During the annual budget review
- C Only after the pilot ends
- D In the first week before the pilot runs

7. What role do acceptance criteria play in a 30-day AI pilot?

- A Scheduling the parallel run's start date
- B Deciding in advance what results count as success
- C Choosing which employee owns the workflow
- D Tracking how much the subscription costs

8. What turns an AI experiment into a dependable business capability?

- A. A documented workflow that someone owns and reviews
- B. Removing people from every decision
- C. A famous prompt
- D. More subscriptions

9. Which item does NOT belong in a short one-page AI use policy?

- A. External-use rules and an owner
- B. A guaranteed monthly revenue target
- C. Review rules and incident reporting
- D. Approved tools and prohibited data

10. What is the purpose of a kill switch in an AI operating system?

- A. To log every action for accountability
- B. To require human approval before actions
- C. To rapidly shut the workflow down
- D. To store tested prompts and examples

11. An escalation rule tells the workflow what to do in which situation?

- A. When outputs are routine and low-risk
- B. When a prompt needs to be rewritten
- C. When the subscription is due for renewal
- D. When uncertainty or an exception arises

12. Which control in the operating system exists primarily for accountability?

- A. The audit log
- B. The prompt library
- C. The kill switch
- D. The approval gate

Go deeper

ONLINE This lesson on the web flashfeedlearningapp.com/study-guides/ai-for-small-business/a-30-day-implementation-playbook

STUDY AI for Small Business in the FlashFeed app flashfeedlearningapp.com/packs/ai-for-small-business

Answer key

Lesson 1: AI Foundations for an Owner

1. C Verify each statistic against reliable sources

Fluency is not evidence. Verify material claims using reliable sources before using them.

2. B A confident-sounding output that contains invented or unsupported information

A hallucination is a confident-sounding AI output containing invented, unsupported, or incorrect information. The other options describe unrelated behaviors, not fabricated content.

3. A Finding patterns in a batch of customer reviews you upload

Analysis finds patterns or implications in material you provide, like uploaded reviews. Drafting, writing, and brainstorming all create new drafts, which is generation.

4. D Fluent, confident wording confirms the claim is accurate

Fluency is not evidence of accuracy, so that statement is false. A polished answer can be wrong or stale, outputs vary, and claims should be verified.

5. B They used different populations, definitions, and methods

The figures differ because they surveyed different populations with different definitions and methods. Always ask who was surveyed and what counted as use before repeating a headline.

6. A The population surveyed and how 'use' was defined

Adoption figures vary by population and definition, so check who was surveyed and what counted as use before repeating a number. Repetition, reputation, or a chatbot's agreement does not settle what was measured.

7. C Verify it against an independent source first

A precise-sounding figure without a traceable source is a warning sign; material claims must be independently verified before use.

8. D Draft weekly social posts from approved offers

Drafting posts from approved offers is frequent, easy to review, and reversible. Auto-sent complaint replies skip review, a yearly plan is infrequent, and a tool comparison is not a workflow.

9. D A repeatable sequence of inputs, decisions, actions, and outputs

A workflow is a repeatable sequence of inputs, decisions, actions, and outputs used to complete a business task. The other options describe purchases or costs, not the process itself.

10. B Deciding whether to fire an employee

Firing an employee has ambiguous inputs, legal consequences, and an irreversible decision, making it a poor first pilot. The others are frequent, reviewable, and low consequence.

11. A It exposes errors, exceptions, and missing instructions

A parallel run reveals errors, exceptions, and missing instructions that a generic demo never shows. It means comparing time, quality, and correction effort, and the owner still defines what good output looks like.

12. C Frequent, easy to check, and low consequence if wrong

A good first task is frequent, structured, easy to review, and low consequence when the first draft is wrong. The other traits describe high-risk, poor-fit candidates.

Lesson 2: Prompting as Business Communication

1. B Draft a 120-word reply to a delayed-order complaint using the policy below; acknowledge the delay, avoid promising a date, and offer two next steps

The strongest request supplies task, context, constraints, and a concrete output target.

2. A Task, context, audience, constraints, and expected output

A good prompt is a short business brief with the task, relevant context, the audience, constraints, and the expected output.

3. C Acceptance criteria

Acceptance criteria are the specific conditions an output must meet to be usable; a prompt library and draft pass are separate concepts.

4. D It shows rhythm, warmth, terminology, and level of detail better than adjectives

Examples communicate voice and detail more concretely than adjectives. They do not remove the need for review, and sensitive details should be stripped first.

5. D The model's temperature and random seed

Criteria cover required facts, length, format, prohibited claims, and policies. Model temperature and seed are technical settings, not acceptance criteria.

6. B The expected output format

A bad structure may mean the expected format was undefined. The facts were right, so context is not the gap; naming the format, such as a table, fixes the shape.

7. C Add specific business context and a good example, then revise

Generic output usually signals missing context or examples, not necessarily the wrong tool.

8. A Prompt library

A prompt library stores tested instructions, examples, and checklists for reuse. The other terms name single steps or conditions, not the stored collection.

9. B Identify weaknesses against explicit criteria

The critique pass checks the draft against explicit criteria for weaknesses and unsupported claims. Creating the first workable version is the draft pass.

10. A The instructions, source material, examples, checklist, and owner

Saving the full pattern (instructions, sources, examples, checklist, owner, and date) is more reusable than keeping just one finished output.

11. C The constraints

An unsafe promise signals incomplete constraints. A wrong fact points to missing sources, generic language to vague examples, and bad structure to an undefined format.

12. D Instructions, sources, examples, and checklist together

A reusable prompt-library asset saves the full package -- instructions, sources, examples, and checklist -- not just one finished output.

13. C Dating entries and retesting when policy or model changes

Dating entries and retesting them when a model or policy changes keeps a prompt-library asset trustworthy over time.

Lesson 3: Research, Data, and Decisions

1. B Current local sales, search, survey, or government data with a defined comparison period

Decision-quality evidence needs a source, definition, location, and comparison period.

2. C An original customer survey your business conducted

A primary source is original data or firsthand evidence. Your own survey is firsthand; blogs, news articles, and marketing pages are secondary summaries or positioning.

3. D Comparing customer interviews, transaction data, and public statistics

Triangulation checks a claim against multiple independent methods. Articles repeating one study, or one model's own guesses, are not independent.

4. A The meaning of a standard arithmetic operation

Current search matters when the answer could have changed: prices, rates, and features. A fixed arithmetic definition is stable and does not need fresh sources.

5. D The surveys asked different populations and defined using AI differently

Estimates diverge because surveys sample different populations and define 'using AI' differently, so the definition changes the number.

6. B Cross-checking sales data, interviews, and a government survey

Triangulation means checking a claim against multiple independent sources or methods, such as sales data, interviews, and public statistics.

7. C Repeated articles often share one unverified source

Articles that repeat one another usually trace back to a single unverified source, so their agreement is not independent confirmation.

8. A Inspect data quality and clarify column meanings

Bad or misunderstood input can produce convincing but useless analysis.

9. C The new supplier caused the increase in returns

An observation is directly supported by the data; an inference is an explanation needing more evidence. 'The supplier caused it' interprets the cause and requires proof.

10. B Tests how a result changes when one important input changes

Sensitivity analysis changes one important assumption at a time to show whether a recommendation is robust or fragile.

11. A It reveals which assumptions drive the decision

Scenarios expose the assumptions carrying the result and show where a contingency plan is needed; they do not guarantee accuracy or replace data checks.

12. D Choosing the chart color scheme

Data-quality inspection covers duplicates, missing values, unclear units, and column meanings. A chart's colors are presentation, not data quality.

Lesson 4: Marketing, Sales, and Customer Service

1. A Remove or substantiate the claim and review it for legal and brand risk

Marketing claims need truthful support. AI does not make an unsupported promise safe.

2. C A month-by-month sales revenue forecast

A brand brief names the customer, promise, tone, proof points, preferred and prohibited terms, and examples. A revenue forecast is a financial projection, not part of the brief.

3. B Adapt it into an email, social posts, a sales handout, and video scripts, each tailored to its channel

A content atom is one substantial source asset adapted into multiple channel-specific pieces. Adaptation matters because each channel serves a different context.

4. D When a human contributes sufficient expressive authorship

The Copyright Office says human expressive authorship can make AI-assisted work protectable, but prompts alone generally do not establish authorship over a purely generated output.

5. C Specific evidence that supports a marketing claim

A proof point is concrete evidence backing a claim. The other options describe prohibited terms, brand voice, and an adapted content atom.

6. D Say it is unsure and escalate to a person

An explicit uncertainty response and handoff protects trust better than a plausible invention.

7. A A defined condition that stops automation and transfers a case to human review

An escalation rule tells the workflow when to stop and hand a case to a person. Ranking prospects by how likely they are to buy describes lead scoring, not escalation.

8. B A routine question about current store hours

Threats, refunds above a limit, and sensitive information are common triggers to route to a person. A routine hours question is exactly what automation should answer.

9. D The business's approved policies, pricing, hours, and canonical answers

Customer-facing AI must be grounded in approved business records. General model knowledge is not a substitute for your own policies and pricing.

10. C Referencing the order the customer just placed with you

Good personalization uses information the customer knowingly supplied for the relationship. The other options reveal hidden or unrelated data the customer did not expect you to use.

11. B Approved policies, product details, hours, and pricing rules

Customer-facing answers should come from the business's own approved policies and records, not general model knowledge.

12. A The business's own approved pricing rules and records

An approved-knowledge design answers from the business's own current pricing rules and records, not from general or guessed information.

Lesson 5: Operations, Automation, and Agents

1. A The actual process demonstrated by an experienced worker

AI can organize knowledge, but the source should be the real process and people who understand it.

2. B Instructions for unusual conditions or failures

The exception path covers unusual conditions or failures (blank field, objection, system down). The normal sequence is the happy path.

3. D A guarantee that no experienced worker input is ever needed

The experienced worker, not the model, is the authority, so their input is essential. Good procedures do include steps, exceptions, an owner, and a review date.

4. C Mark each uncertain rule with a question and resolve it with the process owner

Mark every uncertain rule with a question instead of filling the gap, then resolve it with the process owner. Guessing hides the real expertise.

5. B Test it with a different person and watch where they pause or guess

Testing with a different person reveals where they pause, guess, or need undocumented access. The author's reread or timing, or an AI check, cannot show where a newcomer gets stuck.

6. C Sending a refund from the business account

Moving money is consequential and should stay behind a human approval gate.

7. A It can choose steps and use tools toward a goal

An agent selects steps and uses tools toward a goal, while rule-based automation runs a predefined trigger and action. That flexibility is why agents need firmer boundaries.

8. D A required human confirmation before a consequential action

An approval gate is a required human confirmation before consequential actions like payments or external messages. Logs and least privilege are separate controls.

9. D Read and draft permissions

Least privilege means beginning with read and draft permissions and expanding only when measured performance justifies it. Send, spend, and delete are the riskiest.

10. C Copying new form submissions into a CRM

Rule-based automation reliably handles fixed triggers like copying form submissions. AI earns its place only when a workflow must interpret variable text, classify exceptions, summarize, or draft.

Lesson 6: High-Stakes Work: Money, Contracts, and People

1. D Use the flag to prepare questions for a qualified attorney

AI can focus professional review, but the legal judgment and negotiation should come from a qualified person.

2. B A provision important enough to affect cost, rights, obligations, risk, or the decision to sign

A material term is significant enough to influence cost, rights, obligations, risk, or whether to sign. The others describe boilerplate, paper stock, or trivial details, which are not material.

3. C A plausible-looking dashboard total proves the accounts actually agree

A plausible total is not a reconciled total; an attractive dashboard is not proof records agree. The other three reflect a reconcile-to-source approach.

4. A Current law, jurisdiction, structure, elections, and records all matter

General AI may omit elections, exceptions, jurisdictional rules, or recent changes; current law, structure, and records matter. Use AI to prepare records and questions, not to manufacture certainty.

5. A Summarizing a lease and drafting questions for counsel

AI's first pass is to summarize, classify, compare, or prepare questions. The final filing, commitment, payment, or approval stays with a qualified, responsible person.

6. B Draft a job description from documented role requirements for human review

Drafting from documented requirements is administrative and reviewable; automated judgments about people carry much greater risk.

7. D A substantially different negative outcome for a protected group from a seemingly neutral practice

Adverse impact is a substantially different negative outcome for a protected group caused by a seemingly neutral practice. It need not be intentional, which distinguishes it from an openly discriminatory policy.

8. C Screen out people with disabilities

The EEOC and DOJ warn such tools may screen out people with disabilities, so employers should offer an accommodation process and an alternative. Cost, speed, and data exposure were not the subject of that warning.

9. A Authority, time, relevant information, and ability to challenge the recommendation

Meaningful review is more than clicking approve; the reviewer needs authority, time, relevant information, and the ability to challenge the result, then document the criteria and decision.

10. D Infer a candidate's health, age, or family status

Do not ask AI to infer personality, health, family status, age, race, religion, or other protected traits. The other tasks are administrative uses tied to documented, job-related criteria.

11. C To detect a disparate negative outcome for a protected group

A neutral-looking process can still produce adverse impact, so outcomes must be checked for a substantially different negative effect on a protected group.

12. B Screening out qualified people with disabilities

The EEOC and DOJ have warned that algorithmic tools may screen out people with disabilities, a form of adverse impact requiring an accommodation process.

Lesson 7: Privacy, Security, and Responsible Use

1. D Verify authorization, tool controls, needed fields, and retention before sharing data

Sensitive-data work starts with purpose, authorization, minimization, vendor controls, and retention.

2. C Share only the fields actually needed for the task

Data minimization means sharing only the fields needed for the task, which reduces exposure and often improves focus. Sharing everything raises risk.

3. B They may differ, so verify the current terms for that exact product

Training, retention, and sharing can differ by product and account type, so you must verify the current controls for the exact product. Plan labels alone are not data controls.

4. A A marketing brochure published on your website

Public data is information intentionally available to anyone, like published marketing copy. Customer lists, payroll files, and confidential contracts are restricted business information.

5. D How long a service or business keeps information before deletion or archival

Data retention is how long information is kept before deletion or archival. Access counts and training use are separate vendor questions.

6. B Share only the fields needed for the task

Data minimization means sharing only what a task requires, which reduces exposure of unnecessary information.

7. C Remove fields the task does not need

Minimization means removing fields beyond what the task requires, not just swapping out direct identifiers.

8. A Default training use and data retention period

Vendor evaluation should cover default training use and data retention period before sensitive data is shared.

9. B Business, Enterprise, and API

As of 2026, OpenAI's Business, Enterprise, and API offerings do not use inputs and outputs for training by default — verify the current policy, since vendor defaults can change.

10. D Automatically changing employee pay based on an AI score

Pay decisions affect rights and livelihoods, so they require strong controls and accountable human judgment.

11. C An action that can materially affect money, rights, safety, or trust

A consequential action can materially affect money, rights, safety, or trust, so it needs stronger controls. An action correctable with limited harm is a reversible one, not a consequential one.

12. A At a point where review can still change the outcome

A human in the loop must sit where review can still change the outcome, with approval gates before consequential actions. Review after an irreversible action cannot prevent harm.

13. A Deleting all logs to reduce liability

An incident plan says to preserve records, not delete them. Stopping the workflow, correcting affected data, and notifying owners are all core steps.

14. D Low risk

Internal, low-stakes, easily undone work is low risk. Medium reaches customers, high affects money or rights, and prohibited uses are those whose risk cannot be managed.

Lesson 8: A 30-Day Implementation Playbook

1. D When measured results meet criteria and risks are controlled

Expansion should follow evidence and control, not novelty or sunk cost.

2. C A time-boxed implementation run under clear ownership and success criteria

A pilot is a limited, controlled implementation that exists to create evidence, not to prove AI is exciting or to launch everywhere immediately.

3. A The final expansion decision

A baseline captures current time, cost, error rate, turnaround, backlog, and customer outcome. The expansion decision comes at the end, after comparing results to that baseline.

4. B By running a parallel test beside the current process

A parallel run works the AI beside the existing process so you can record where it succeeds, fails, and how long review takes without risking live work.

5. C A frequent, structured, reviewable, low-risk workflow

The pilot should target a frequent, structured, reviewable, low-risk workflow so results are measurable and mistakes are recoverable.

6. D In the first week before the pilot runs

Acceptance criteria should be defined in the first week, alongside the baseline and the data the pilot may use.

7. B Deciding in advance what results count as success

Acceptance criteria set the standard results must meet, so the later expansion decision is evidence-based, not a guess.

8. A A documented workflow that someone owns and reviews

Dependability comes from the operating system around the model, not the model alone.

9. B A guaranteed monthly revenue target

A beginner AI policy covers approved tools, prohibited data, review rules, external use, incidents, and responsibility. A revenue target is not part of a use policy.

10. C To rapidly shut the workflow down

A kill switch exists for rapid shutdown. The audit log handles accountability, the prompt library stores prompts, and the approval gate handles human sign-off.

11. D When uncertainty or an exception arises

An escalation rule routes uncertainty and exceptions to a person, rather than letting the AI proceed unchecked on cases it cannot handle.

12. A The audit log

An audit log records what happened for accountability. The approval gate blocks consequential actions, the library stores prompts, and the kill switch enables shutdown.

Sources

Lesson sources: U.S. Small Business Administration — AI for small business; U.S. Census Bureau — AI use by U.S. businesses (May 2026); U.S. Chamber of Commerce — 2025 Empowering Small Business report; Ethan Mollick — Using AI Right Now: A Quick Guide; OpenAI Academy — Small Business Resource Hub; U.S. Copyright Office — Copyright and AI, Part 2; NIST — Generative AI Profile (NIST AI 600-1); EEOC and DOJ — AI and disability discrimination in employment; NIST — AI Risk Management Framework; Federal Trade Commission — Privacy and security guidance; OpenAI — Business data privacy, security, and compliance; OpenAI Help Center — Data controls in ChatGPT

- U.S. Small Business Administration — AI for small business — https://www.sba.gov/business-guide/manage-your-business/ai-small-business
- U.S. Census Bureau — AI use by U.S. businesses (May 2026) — https://www.census.gov/library/stories/2026/05/ai-use-businesses.html
- U.S. Chamber of Commerce — 2025 Empowering Small Business report — https://www.uschamber.com/technology/artificial-intelligence/u-s-chambers-latest-empowering-small-business-report-shows-majority-of-businesses-in-all-50-states-are-embracing-ai
- Ethan Mollick — Using AI Right Now: A Quick Guide — https://www.oneusefulthing.org/p/using-ai-right-now-a-quick-guide
- OpenAI Academy — Small Business Resource Hub — https://academy.openai.com/en/public/resources/openai-academy-small-business-resource-hub-2026-06-03
- U.S. Copyright Office — Copyright and AI, Part 2 — https://www.copyright.gov/ai/Copyright-and-Artificial-Intelligence-Part-2-Copyrightability-Report.pdf
- NIST — Generative AI Profile (NIST AI 600-1) — https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf
- EEOC and DOJ — AI and disability discrimination in employment — https://www.eeoc.gov/newsroom/us-eeoc-and-us-department-justice-warn-against-disability-discrimination
- NIST — AI Risk Management Framework — https://airc.nist.gov/airmf-resources/airmf/
- Federal Trade Commission — Privacy and security guidance — https://www.ftc.gov/business-guidance/privacy-security
- OpenAI — Business data privacy, security, and compliance — https://openai.com/business-data/
- OpenAI Help Center — Data controls in ChatGPT — https://help.openai.com/en/articles/7730893-data-controls-faq
- OpenAI — Enterprise privacy — https://openai.com/enterprise-privacy/

Image credits

Each picture in this guide is credited below, with its creator, source and license where there is one.

- **Lesson 1, A Pattern Machine, Not an Oracle:** FlashFeed original diagram · CC-BY-4.0
- **Lesson 2, Prompt Like a Brief:** FlashFeed original diagram · CC-BY-4.0
- **Lesson 3, Current Questions Need Current Sources:** FlashFeed original diagram · CC-BY-4.0
- **Lesson 3, Use Scenarios, Not One Forecast:** Wikimedia Commons · Shixart1985 · CC BY 2.0 — https://commons.wikimedia.org/wiki/File:Agreeing_on_a_business_plan_at_home.jpg
- **Lesson 5, Automation Starts with a Trigger:** FlashFeed original diagram · CC-BY-4.0
- **Lesson 7, Share the Minimum:** AI-generated (gpt-image-1.5)
- **Lesson 7, Control by Consequence:** AI-generated (gpt-image-1.5)
- **Lesson 8, Keep the First Policy Short:** AI-generated (gpt-image-1.5)

Useful links

FlashFeed

- This pack in the app: flashfeedlearningapp.com/packs/ai-for-small-business
- All packs: flashfeedlearningapp.com/packs
- How FlashFeed works: flashfeedlearningapp.com/how-it-works
- Free and open: flashfeedlearningapp.com/free-and-open

About FlashFeed

Learn anything, a little at a time.

FlashFeed turns big subjects into Knowledge Packs of short cards, quick games and study guides. Scroll when you have a minute. Go deeper when you have more. Free, with no ads and no account to make.

flashfeedlearningapp.com

Download on the App Store

Get it on Google Play

Follow

- Instagram — @flashfeedlearningapp · @flashfeed_history (history) · @flashfeed_spanish (Spanish)
- YouTube — @FlashFeedLearningApp
- TikTok — @flashfeed.learningapp
- FlashFeed Study Notes podcast — Spotify · Apple Podcasts · RSS

FlashFeed, by Recurxive LLC. November 2026 edition, built from AI for Small Business pack version 0.1.18. Typeset in Big Shoulders Display and Figtree (SIL Open Font License). Pictures are credited on the image credits page.