PUBLIC PREVIEW

Small business AI buying kit

A two page sample of the buyer briefing

PREVIEW EDITION | Full package not yet on sale

Two-page public preview. Full package not yet on sale. Prepared 31 October 2026.

Before buying an AI tool, define one workflow and test the complete job: the input, the AI output, human review and what happens next. A fast first draft is useful only when the final result is good enough and the data access is appropriate.

Three questions before a subscription

Question	What to establish
What work will improve?	One repeatable task, a measured manual baseline and an agreed quality standard.
What data can the tool reach?	The exact plan, approved sources, necessary permissions and applicable data-use terms.
What benefit will we actually use?	Complete cash spending, human handling and review time, and a specific use for released capacity.

A first pilot should be easy to reverse

A small business could test drafts of routine customer replies using an approved FAQ. Staff still check the answer and send it. The trial excludes sensitive customer records and gives the tool no permission to send, modify accounts or issue refunds.

Before testing, define acceptable answers, critical errors, the task mix, a cash limit and a stop condition. Keep a manual fallback. Use representative tasks, including awkward questions and missing information; record corrections and review time, not just generation speed.

A claim needs a source and a scope

For each important claim, save the product and plan, exact wording, source or contract clause, date checked and owner. Mark missing evidence as unknown. A general assurance about “business security” does not establish the permissions, retention or export behavior of your chosen configuration.

The full kit provides a buyer briefing, an editable checklist and evidence-request template, a two-week pilot decision worksheet and a local cost calculator. These are decision aids, not vendor certifications or guarantees.

SAMPLE A HYPOTHETICAL COST CHECK

Do not mistake capacity for cash savings

Suppose five users pay an assumed $30 per seat each month, with $40 of usage and $25 for integrations. One-time implementation costs $1,200, plus 20 internal hours valued at $35 an hour. These are invented assumptions, not current vendor prices.

Measure	Hypothetical result
Recurring cash spending	5 × $30 + $40 + $25 = $215 per month
First-year cash spending	12 × $215 + $1,200 = $3,780
Manual workload	400 tasks × 8 minutes = 53.33 hours per month
AI handling and review	400 × (3 + 2) minutes = 33.33 hours per month
Time released	20 hours per month; $700 of capacity at $35 per hour
Proven recurring cash savings	$0 per month
First-year economic cost	$18,480 including cash, setup time, handling and review

The 20 hours may be valuable if the team can use them for other work. They do not reduce the cash bill by themselves. The $700 capacity value is another way to express those same hours; never count both as separate benefits or call that amount cash saved.

The economic total adds $3,780 cash, $700 internal setup time and $14,000 of annual handling and review labor value. Review costs $466.67 per month and handling costs $700 per month. Do not add the same labor twice or deduct hypothetical capacity from cash spending.

Check the control you are relying on

Retention depends on the product, feature and agreement. Permission-respecting search can still surface information that was already overshared. A spending alert may not stop usage. Obtain plan-specific evidence and test the control before relying on it. [2, 3, 4, 5]

Source examples

[2] Anthropic commercial data retention policy

[3] Claude API and data retention

[4] Microsoft Copilot data protection and auditing

[5] OpenAI API usage and spend limits

Source examples checked 31 October 2026. Recheck before purchase. The source descriptions illustrate questions to ask; they do not rank or endorse providers. The calculations and pilot approach are original working methods.