

Business Agent Launch Pack

Build one narrow reusable agent with repository guidance, a custom skill, explicit permissions and human gates, real and held-back evals, a passing run record, and exact launch, update, disable, and rollback instructions.

Course: <https://flowgrammer.ca/courses/build-business-ai-agent>

Complete one section after each lesson. Use permitted real or sanitized evidence when available. Keep fictional fixtures, unknown values, manual work, blocked routes, and human approvals clearly labelled.

1. Define One Agent Job and Outcome

A job brief, v0 boundary, and three to six independent pass/fail outcomes.

Begin with the business decision and finished artifact, not with a model or agent label. A narrow job is easier to test, own, repair, and stop.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] The job can be completed end to end without a second workflow. [] Every outcome has observable evidence. [] External effects are excluded or explicitly gated. [] A held-back input is reserved.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

2. Choose the Smallest Native Surface

A surface decision that explains why repository guidance plus one skill fits the chosen job.

Different agent artifacts solve different problems. Choose the smallest surface that makes the job durable without turning a simple workflow into infrastructure.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

☐ The chosen surface matches the recurring need. ☐ No separate role exists without a distinct boundary. ☐ Manual operation remains possible. ☐ Deferred features have mechanisms rather than vague wishes.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

3. Write Durable Repository Guidance

A scoped AGENTS.md section that a fresh run can use without loading every manual.

Repository guidance should change decisions across relevant work. It should not duplicate the entire skill or accumulate every edge case ever observed.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] A fresh reader can identify authoritative context. [] The file routes rather than duplicates detailed procedures. [] Unknown and blocked states are permitted. [] No secret or private record appears.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

4. Build a Concise Custom Skill

A validated skill folder with clear activation, minimal resources, and no unsafe default.

A useful skill supplies non-obvious workflow decisions without constraining unrelated work. Its discovery text must be precise before the full instructions load.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] The description is discriminating. [] Conditional references are directly discoverable. [] Scripts have declared inputs and failures. [] The skill has no secret or external write default.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

5. Set Permissions and Human Gates

A permission matrix and stop policy tied to the actual job and tools.

Instructions should make authority visible, but effective control comes from the runtime, tools, and approval path. Test the boundary you claim.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Every consequential action has a trusted gate. [] The runtime limitation is disclosed. [] Source text cannot change authority. [] The human owns policy and relationship decisions.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

6. Create Real and Held-Back Evals

An eval suite with prewritten expectations, labelled fixtures, and evidence requirements.

An eval describes what should happen before the agent produces an answer. It protects the outcome from being rewritten to match a convenient result.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Fiction and real evidence are labelled separately. [] Expected results predate observed outputs. [] Boundary tests cause no real side effects. [] Every verdict points to evidence.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

7. Run, Grade, Repair, and Record

A passing run log or an honest blocked result after focused repair cycles.

A polished answer is not a verdict. Grade the declared outcomes, identify the failing layer, and change one thing at a time.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Failed evidence remains available. [] Expected results do not move to match output. [] The held-back case passes independently. [] No untested criterion is marked pass.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

8. Verify Launch, Rollback, and Maintenance

A complete Business Agent Launch Pack and a final launch, manual, blocked, or stop decision.

The capstone is complete only when another run can discover and use the installed artifact, and an operator can turn it off or restore the prior version.

Decision, scope, and evidence used

Item or rule	Evidence / source	Unknown or failure action	Owner

[] Installed behaviour matches staged behaviour. [] Rollback is executable and names the prior known-good artifact. [] No secret entered instructions or logs. [] Schedule and external integrations remain unclaimed unless tested.

Result: pass, repair, manual, blocked, narrower scope, or stop, with the reason

9. Capstone and critical gates

Grade the complete Business Agent Launch Pack and choose the next operating decision.

Criterion	Missing	Partial	Decision-ready	Score

Criteria: Narrow job and binary outcomes; Appropriate native surface; Durable repository guidance; Validated custom skill; Permission and human-gate matrix; Real and held-back evals; Passing run evidence; Installed-copy proof; Launch and rollback instructions.

Critical gates: no unsupported claim, no hidden external effect, no invented missing value, human decisions remain owned, and the tested status matches the evidence. Passing standard: at least 16 of 18 points with no missing critical gate.

Final decision, approver, limitations, rollback or stop route, and next review