

AI Agents for Startup Strategy — Workbook

Len P. van der Hof

2026

AI Agents for Startup Strategy: Workbook

Living templates from *AI Agents for Startup Strategy*. Copy each into your notes system or print it. These are operating instruments, not one-time exercises: every artifact carries a review rhythm. Printable versions: the QR at the back of the book or lenvanderhof.com/ai-agents/worksheets.

Artifact 1 — Information-bottleneck map (Chapter 0)

When: One evening to start; then a one-line entry whenever you are surprised by something you could have known.

Decision (last quarter)	When you decided	When the info became available	When you <i>could</i> have known	What the delay cost
1.				
2.				
3.				

Trace drill: For the worst gap, write the path the signal took to reach you (page → thread → person → meeting) and circle where it sat waiting for a human. That waiting room is your first agent's job description.

Surprise log (standing): date · what surprised you · when it became knowable · when it reached you · cost. A month of entries is your prioritized build list.

Artifact 2 — SENSE design canvas (Chapter 2; extended in 5, 8, 12)

One strategic question: _____

Phase	Your one-line design	Failure mode to avoid
Scan	what is watched, and why those sources	thesis-free collection
Evaluate	the claim shape, against which named exposure	summary, not judgment

Phase	Your one-line design	Failure mode to avoid
Navigate	what may change without asking	the static watchlist
Simulate	the two futures kept warm	confident fortune-telling
Execute	which existing meeting; who owns the output	the digest that arrives nowhere

Cadence ring it feeds: daily / weekly / monthly / quarterly / annual **Kill criterion (write it now):**

Artifact 3 — Delegation-boundary worksheet (Chapter 2; risk modifiers Chapter 10)

Decision	Lane (delegate / augment / reserve)	One-line rationale
1.		
2.		
3.		
4.		
5.		
6.		
7.		
8.		
9.		

Risk re-check (quarterly): Which lanes did hallucination, bias, over-reliance, or regulatory exposure shift this quarter? Move the rows; date the change.

Artifact 4 — Signal map (Chapter 3; assumption register Chapter 4; customer streams Chapter 6)

Three competitive fears, as full sentences:

1. _____
2. _____
3. _____

Watched surface	Quadrant (core / selective / ambient / ignore)	Latency budget

Ignore list (at least three, with reasons): _____

Assumption register (market edition):

#	Assumption (falsifiable sentence)	Tier-1 confirmer	Tier-2 interpreter	Tier-3 early rumble	Timestamp
1					
2					
3					
4					

Customer-signal streams: support · sales notes · churn notes · usage · reviews — mark which currently reach roadmap decisions, and which sit unread.

Artifact 5 — Integration cadence plan (Chapters 7, 9, 13)

System: _____ Existing meeting its output enters: _____ Who would notice if it stopped: _____

Output	Integration pattern (parallel / sequential / human-gate / autonomous+audit)	Gate placement

Load test: the chore this system removes on the user’s first interaction: _____

Investor-graph instance (Ch 7): five most-likely leads, each with thesis (quoted), stage fit, timing signal, warm path. GTM loop instance (Ch 13): standing goal · SENSE walk · baseline metric and current value · guardrails (grounding source, voice reference, volume limit, compliance gates).

Artifact 6 — Five-risk guardrail audit (Chapter 10)

Risk	How it would manifest <i>here</i>	What stands between failure and a decision today	Guardrail to add (class)
Hallucination			
Bias			
Over-reliance			
Security (prompt injection)			
Regulatory			

Evidence-trace test: pick one real claim the system produced and walk it back to its source. If you cannot, the trace is broken — fix that before anything else. **Lane shifts:** which delegate-lane decisions did this audit push toward augment or reserve?

Artifact 7 — Agent-fluency scorecard (Chapter 11)

Score 1–5, quarterly. Lowest score is the quarter’s development priority.

Role	Deployment	Judgment	Correction	Governance
Founder				
System owner(s)				
Fluency lead				
Team average				

Calibration meeting (monthly, 30 min): What did our systems get right that we acted on, and what did it teach us about trust? What did they get wrong, and what was the root cause? What changes, and who owns it?

System metrics: decision influence · claim accuracy · correction-rate trend · time-to-trust-recovery.

Artifact 8 — Override doctrine + 90-day plan (Chapters 14 and Conclusion)

Never delegated (with one line each on why it is constitutively yours):

1. _____
2. _____
3. _____

Override conditions (when I back my judgment over a confident system): missing context I hold · stakes beyond any track record · conflict with the company we are building.

Sharpness practice: three decisions this month my systems could make that I will make myself: 1. _____ 2. _____ 3. _____

The 90-day plan

Days 1–30 · Diagnose (build nothing): run Artifact 1 for a full month; write one Artifact 2 canvas and one Artifact 3 sheet for the single system the log justifies. **Gate:** say it in one sentence — *watches X, claims Y, feeds decision Z in meeting W.*

Days 31–60 · Deploy: one system, simplest architecture, into the existing meeting, instrumented from day one. **Gate:** it changed at least one real decision, and you can point to which.

Days 61–90 · Evaluate: score against the kill criterion you wrote in week one; decide *expand / adjust / kill*, in writing.

The day-30 gate goes in your calendar now: ____ / ____ / ____

The layer is not what you bought. It is what you run.