


ORIA FIELD GUIDE | AUGUST 2026

THE STRATEGY AGENT PLAYBOOK FOR CLAUDE

From an open question to an evidence-backed decision
and an executive-ready deck.

20 bounded agents | DIAGNOSE → ANALYZE → DECIDE → EXECUTE → COMMUNICATE



Claude becomes useful in consequential work when prompts become governed work orders, repeatable methods become Skills, and narrowly scoped agents hand checked analysis to human reviewers.

ORIA

AI for complex slides - after the analysis is approved.

01 | OPERATING SYSTEM

Build the system, not a longer prompt.

A strong prompt improves one answer. A governed operating model improves the next hundred assignments - while keeping authority, evidence, versions, and review visible.

PROJECT CONTEXT	Mandate, definitions, policies, source hierarchy, templates, and disclosure rules that remain stable across runs.
WORK ORDER	This run's objective, evidence, method, deliverables, permissions, stop conditions, review gates, and acceptance tests.
SKILL	A versioned method for work you repeat: source ledger, model audit, scenario design, storyline, or deck QC.
AGENT	A bounded workflow owner with instructions, permitted tools, an output contract, and a named handoff - never a decision-maker.
CONNECTOR / MCP	Governed access to approved current data or systems. Start read-only; access is not proof that data is correct or entitled.

SET UP ONE AGENT BEFORE YOU BUILD THE LIBRARY

- Choose one recurring decision and define the artifact, evidence, tie-outs, reviewer, and prohibited actions.
- Give the agent the smallest tool and Skill surface it needs; stop on missing evidence, conflicts, or permission ambiguity.
- Test the standard case, a conflicting-input case, and a case that must escalate. Version the definition and record the owner.

OFFICIAL SOURCES

[Anthropic: effective agents](#)
[Claude Code: subagents](#)
[Agent Skills overview](#)

02 | PROMPTS + SKILLS

Write the work order. Record the repeatable method.

The best strategy prompt is inspectable: it makes the decision, source boundary, method, authority, and definition of done explicit before Claude calculates or writes.

01

OBJECTIVE + DECISION

Question, owner, deadline, and action the work may support.

02

DELIVERABLES

Exact model, memo, storyboard, tables, schema, and file locations.

03

EVIDENCE

Approved sources, versions, as-of date, permissions, and conflict hierarchy.

04

METHOD + CONVENTIONS

Calculations, scenarios, units, signs, periods, accounting basis, and materiality.

05

AUTHORITY + GATES

What Claude may change, prohibited actions, stops, reviewers, and approvers.

06

ACCEPTANCE TESTS

Tie-outs, tolerances, reproducibility, citations, and cross-artifact consistency.

COPY-READY FIRST RESPONSE

Restate the objective, decision use, approved evidence, conventions, deliverables, permissions, stops, and acceptance tests. List missing inputs and conflicts. Propose the smallest agent route and review gates. Return the source ledger before analysis. Label material claims FACT, CALCULATION, ASSUMPTION, or HYPOTHESIS. Do not calculate, edit, publish, contact anyone, or take an external action until the named reviewer approves the route.

RECORD A SKILL WHEN THE METHOD SURVIVES REVIEW

- Capture why each source is trusted, how periods and units are normalized, where judgment enters, which exceptions stop the run, and what done means.
- Keep the stable procedure in SKILL.md; put detailed references, deterministic checks, and templates in references/, scripts/, and assets/ only when needed.
- Test normal, edge, stale-source, negative-trigger, and approval-required cases. Remove secrets and sensitive examples before recording or sharing.
- For slides, record Orias audience, action-title, exhibit, sourcing, template, and deck-QC method - then pass only approved claims into production.

OFFICIAL SOURCES

Prompting best practices

Skill authoring

Create custom Skills

03 | AGENT DIRECTORY | DIAGNOSE + ANALYZE

Use the smallest specialist set that can finish the work.

01-05 of 20 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

01

Name: PROBLEM FRAMER

What it does: Turns a topic into one decision with owner, deadline, objective function, SCQ, constraints, and success criteria.

When to use: The brief says “build an AI strategy” or different leaders are answering different questions.

[Download the companion agent pack](#)

02

Name: SITUATION ASSESSOR

What it does: Triangulates financial, market, customer, and operational evidence into a current-state baseline.

When to use: Leadership disagrees on what is happening or is debating solutions before diagnosing causes.

[Download the companion agent pack](#)

03

Name: HYPOTHESIS BUILDER

What it does: Converts beliefs into falsifiable hypotheses, must-be-true assumptions, and kill-shot tests.

When to use: The team “already knows the answer,” evidence collection is sprawling, or confirmation bias is likely.

[Download the companion agent pack](#)

04

Name: ISSUE TREE ARCHITECT

What it does: Decomposes the approved decision into answerable branches and prioritizes the critical path.

When to use: The workplan is broad, duplicated, or disconnected from the decision.

[Download the companion agent pack](#)

05

Name: MARKET SIZER

What it does: Defines boundaries and triangulates top-down and bottom-up market ranges, growth, and serviceable segments.

When to use: Entry, adjacency, product, or portfolio choices depend on market scale and growth.

[Download the companion agent pack](#)

Use the smallest specialist set that can finish the work.

06-10 of 20 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

06

Name: PROFIT POOL MAPPER

What it does: Maps revenue, margin, and value capture across the chain and tests where profit is migrating.

When to use: Revenue scale may conceal unattractive economics or value is shifting between ecosystem roles.

[Download the companion agent pack](#)

07

Name: COMPETITOR WAR-GAMER

What it does: Builds an authorized competitor fact base, infers responses with confidence labels, then simulates moves and counter moves.

When to use: Pricing, launch, entry, or M&A choices will trigger consequential rival or stakeholder reactions.

[Download the companion agent pack](#)

08

Name: CUSTOMER SEGMENTER

What it does: Creates behavior- and needs-based segments from interviews, CRM, support, win/loss, and survey evidence.

When to use: Demand, proposition, adoption, retention, or willingness-to-pay assumptions drive the decision.

[Download the companion agent pack](#)

09

Name: STRATEGIC OPTIONS GRID

What it does: Designs 3-4 coherent choice systems across where to play, how to win, capabilities, management systems, economics, and explicit non-choices.

When to use: A decision is framed as yes/no or token alternatives are disguising a preferred answer.

[Download the companion agent pack](#)

10

Name: PRICING STRATEGIST

What it does: Diagnoses value, structure, and execution; designs pricing options tied to customer evidence and unit economics.

When to use: Growth, margin, packaging, discount leakage, or monetization is central to the decision.

[Download the companion agent pack](#)

05 | AGENT DIRECTORY | DECIDE + EXECUTE

Use the smallest specialist set that can finish the work.

11-15 of 20 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

11

Name: BUSINESS CASE BUILDER

What it does: Translates each option into driver-based P&L, cash, NPV/ROIC/payback logic, sensitivities, and break-even conditions.

When to use: Capital, pricing, share, timing, cost, or value at risk could change the choice.

[Download the companion agent pack](#)

12

Name: PORTFOLIO REVIEWER

What it does: Compares businesses, products, markets, or initiatives on attractiveness, right to win, returns, risk, and strategic role.

When to use: Capital and leadership attention must move across a portfolio rather than into one isolated bet.

[Download the companion agent pack](#)

13

Name: OPERATING MODEL DESIGNER

What it does: Translates the chosen strategy into capabilities, structure, processes, decision rights, governance, talent, data, and technology.

When to use: The choice is clear but the organization is not yet configured to deliver it.

[Download the companion agent pack](#)

14

Name: INITIATIVE PRIORITIZER

What it does: Cuts the initiative list to the few moves capacity can absorb, using value, urgency, feasibility, dependencies, and confidence.

When to use: Everything is labeled priority or delivery capacity is the binding constraint.

[Download the companion agent pack](#)

15

Name: TRANSFORMATION ROADMAPPER

What it does: Sequences approved initiatives into phases with exit criteria, milestones, single accountable owners, and a concrete first 90 days.

When to use: Leadership has approved a strategy but cannot answer what happens Monday.

[Download the companion agent pack](#)

06 | AGENT DIRECTORY | EXECUTE + COMMUNICATE

Use the smallest specialist set that can finish the work.

16-20 of 20 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

16

Name: KPI ARCHITECT

What it does: Links the strategic outcome to leading drivers, lagging results, value realization, action thresholds, owners, and anti-gaming checks.

When to use: A roadmap exists but leaders lack early warning, benefit ownership, or stop/scale rules.

[Download the companion agent pack](#)

17

Name: PYRAMID STORY BUILDER

What it does: Synthesizes approved analysis into one governing answer, 2-4 reasons, evidence, implications, risks, and ask.

When to use: The evidence is complete but the message is chronological, sprawling, or buried.

[Download the companion agent pack](#)

18

Name: DECISION MEMO WRITER

What it does: Produces a concise, answer-first memo with recommendation, rejected alternatives, economics, conditions, dissent, and exact approvals requested.

When to use: Executives need a durable pre-read or decision record rather than a presentation-first artifact.

[Download the companion agent pack](#)

19

Name: HOSTILE Q&A SIMULATOR

What it does: Independently attacks the recommendation from CEO, CFO, operator, board, regulator, customer, and competitor perspectives.

When to use: Stakes are high, dissent is hidden, or the meeting may expose load-bearing assumptions.

[Download the companion agent pack](#)

20

Name: EXECUTIVE DECK PACKAGER

What it does: Converts the approved storyline into a production-ready slide specification with action titles, exhibit briefs, data lineage, source notes, appendix, and speaker cues.

When to use: The recommendation is approved and ready for board, CEO, investment committee, or transformation governance.

[Download the companion agent pack](#)

07 | ROUTING + ROLLOUT

Run the first controlled sprint with a minimum viable team.

The coordinator owns versions, dependencies, approvals, and open items. Specialist agents receive frozen inputs, return structured artifacts, and stop at named human gates.

01	BRIEF	Lock the decision, source boundary, output contract, permissions, and tests.
02	EVIDENCE	Build the source ledger first; resolve conflicts before synthesis.
03	BUILD	Route only independent tasks in parallel; preserve versions and handoffs.
04	CHALLENGE	Use a separate reviewer to reproduce material outputs and attack assumptions.
05	APPROVE	A named professional approves numbers, narrative, confidentiality, and release.
06	ORIA	Hand approved page briefs to Oria for complex, editable slides; then run deck QA.

30-DAY CONTROLLED PILOT

WEEK 1	Select one recurring workflow; freeze sources, baseline quality, reviewers, and failure tests.	WEEK 2	Configure the smallest agent route; keep tools read-only; test conflicts and stops.
WEEK 3	Run live work beside the current process; log corrections, time, provenance, and escalations.	WEEK 4	Approve, revise, or retire the route. Scale only after control and decision quality improve.

THE ORIA PRODUCTION BRIEF

Audience + decision → action-title spine → one message per slide → decisive exhibit + exact source IDs → caveat + review status → appendix proof. Claude structures the approved argument; Oria turns it into beautiful, complex slides.

[Get the companion agent pack](#)[Build complex slides with Oria](#)

OFFICIAL SOURCES

- Managed-agent permissions
- Multiagent orchestration
- Financial services repository