

ORIA FIELD GUIDE | AUGUST 2026

CLAUDE FINANCE AGENT PLAYBOOK

Put Claude to work across planning, operations, close, decisions, and communication.

20 bounded agents | PLAN → OPERATE → CLOSE → DECIDE → 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.

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 finance 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

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: BUDGET ARCHITECT

What it does: Converts strategic objectives into a driver tree, assumption register, funding envelope, planning calendar, and owner map; reconciles bottom-up submissions to top-down targets.

When to use: At the start of the annual budget or a re-plan, after targets and governing actuals are available but before templates are distributed.

[Download the companion agent pack](#)

02

Name: DRIVER-BASED FORECASTER

What it does: Links volume, price, mix, headcount, productivity, commodity, FX, and working-capital drivers to a formula-driven forecast with an assumptions log and checks.

When to use: For rolling forecasts and estimate refreshes when management needs to understand what changed and why.

[Download the companion agent pack](#)

03

Name: SCENARIO MODELER

What it does: Builds internally coherent base, upside, and downside cases; identifies the variables that reverse the conclusion and defines observable triggers for switching cases.

When to use: Before committing to a forecast, refinancing plan, investment, or operating response exposed to material uncertainty.

[Download the companion agent pack](#)

04

Name: CAPITAL ALLOCATION

What it does: Compares reinvestment, acquisition, debt reduction, liquidity reserves, and shareholder returns against approved hurdle rates, constraints, reversibility, and downside protection.

When to use: When the CFO or board must choose among competing uses of cash; finance and the decision owner still approve policy and the final choice.

[Download the companion agent pack](#)

05

Name: CASH FLOW MANAGER

What it does: Reconciles opening to closing cash, builds a short-term liquidity view, explains receipts and disbursements, and surfaces headroom risks and data gaps.

When to use: For weekly cash reviews, a 13-week forecast refresh, or before a funding decision; never to move money or authorize payments.

[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: WORKING CAPITAL OPTIMIZER

What it does: Decomposes DSO, DPO, DIO, overdue balances, inventory, and operational bottlenecks; ranks cash-release levers with owners and trade-offs.

When to use: When cash conversion weakens, liquidity is tight, or an operating review needs targeted receivables, payables, or inventory actions.

[Download the companion agent pack](#)

07

Name: SPEND CONTROLLER

What it does: Compares spend to budget, policy, commitment, and run-rate; classifies exceptions and prepares an owner-routed review queue without approving or rejecting transactions.

When to use: For monthly cost control, procurement review, or early warning on recurring and uncommitted spend.

[Download the companion agent pack](#)

08

Name: REVENUE AND MARGIN MONITOR

What it does: Reconciles reported totals and bridges price, volume, mix, FX, cost, customer, product, and channel effects; separates confirmed drivers from hypotheses.

When to use: In weekly trading, monthly performance, or earnings reviews after the revenue and gross-profit perimeter is locked.

[Download the companion agent pack](#)

09

Name: CLOSE COORDINATOR

What it does: Converts the close calendar and status inputs into dependencies, blockers, evidence status, owners, due dates, and an exception-first daily control view.

When to use: During pre-close and close to coordinate the process; it must not post journals or certify completion.

[Download the companion agent pack](#)

10

Name: RECONCILIATION REVIEWER

What it does: Compares GL, subledger, bank, schedule, and support balances; matches items, quantifies breaks, and produces a reviewer-ready exception log with evidence.

When to use: After preparer work and before controller sign-off, with read-only access and a locked reporting perimeter.

[Download the companion agent pack](#)

05 | AGENT DIRECTORY | CLOSE + DECIDE

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: VARIANCE INVESTIGATOR

What it does: Ranks movements by value, rate, persistence, and materiality; builds defined bridges and labels each explanation fact, calculation, management explanation, or hypothesis.

When to use: For close flux analysis, forecast-versus-actual review, or portfolio performance when causes require evidence rather than plausible prose.

[Download the companion agent pack](#)

12

Name: CONTROLS AND COMPLIANCE MONITOR

What it does: Compares available evidence with the control matrix, flags missing performance or review evidence, and maintains a gap and escalation register.

When to use: For control self-assessment, close readiness, or audit preparation; never to make a legal, audit, accounting, or compliance conclusion.

[Download the companion agent pack](#)

13

Name: FP&A ANALYST

What it does: Connects actual performance to operational drivers, outlook, choices, and actions; builds an answer-first decision brief with assumptions and open questions visible.

When to use: In monthly business reviews, forecast updates, or executive decision support after figures and causes have passed review.

[Download the companion agent pack](#)

14

Name: PROFITABILITY MAPPER

What it does: Maps contribution and fully allocated economics by customer, product, channel, or segment; exposes allocation logic, cost-to-serve, concentration, and cross-subsidy.

When to use: Before portfolio, service-model, pricing, or exit decisions where averages conceal different economics.

[Download the companion agent pack](#)

15

Name: PRICING STRATEGIST

What it does: Models price-volume trade-offs, willingness-to-pay evidence, mix, discount leakage, customer response, margin guardrails, and implementation risks.

When to use: Before a price change, discount-policy revision, packaging decision, or commercial recovery plan; commercial owners approve market assumptions.

[Download the companion agent pack](#)

06 | AGENT DIRECTORY | DECIDE + 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: BUSINESS CASE BUILDER

What it does: Converts options into incremental cash flows, NPV/IRR/payback or approved metrics, sensitivities, failure modes, decision gates, and an explicit recommendation.

When to use: For capex, transformation, product, partnership, or make-buy decisions after tax, accounting, risk, and policy inputs are supplied.

[Download the companion agent pack](#)

17

Name: MANAGEMENT REPORT WRITER

What it does: Converts reviewed performance tables and actions into concise commentary: what changed, why, what it means, what management will do, and which issues remain open.

When to use: For monthly or quarterly management reporting after totals, drivers, and owners have been approved.

[Download the companion agent pack](#)

18

Name: CFO NARRATIVE BUILDER

What it does: Creates an answer-first storyline connecting performance, outlook, risks, choices, and action; maps every claim to evidence and distinguishes certainty from judgment.

When to use: Before a board, investor, lender, or leadership narrative is storyboarded; the accountable executive owns the message.

[Download the companion agent pack](#)

19

Name: BOARD DECK PACKAGER

What it does: Converts an approved storyline, model outputs, and source notes into a slide-by-slide brief with conclusion titles, exhibit specifications, implications, and editable-source mapping.

When to use: Only after the model and narrative pass review. Use Oria to build the final complex slides in the approved template.

[Download the companion agent pack](#)

20

Name: EXECUTIVE Q&A SIMULATOR

What it does: Generates difficult but fair questions, traces each answer to evidence, identifies claims that overreach, and prepares a gap list rather than fabricating a confident response.

When to use: After the draft deck is built and before board, lender, investor, audit committee, or senior-management review.

[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