


ORIA FIELD GUIDE | AUGUST 2026

THE INVESTMENT BANKING GUIDE FOR CLAUDE

From clever prompts to a controlled deal workflow using
21 specialist agents.

21 bounded agents | ORIGINATE → VALUE → MODEL → DILIGENCE → 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 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 Ori's 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 | ORIGINATE + VALUE

Use the smallest specialist set that can finish the work.

01-06 of 21 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

01

Name: SECTOR SCREEN MAPPING

What it does: defines the sector, assembles candidates, applies explicit filters, records provenance, and ranks coverage-not desirability.

When to use: before target selection or pitch prioritization.

[Download the companion agent pack](#)

02

Name: BUYER & ACQUIRER LISTS

What it does: segments strategics and sponsors by documented fit, capacity signals, conflicts, and unknowns; it never contacts them.

When to use: after the asset and screening criteria are approved.

[Download the companion agent pack](#)

03

Name: THESIS & ANGLE DEVELOPMENT

What it does: states a falsifiable thesis, must-be-true assumptions, proof, counter-case, and missing evidence.

When to use: after screening, before pitch narrative.

[Download the companion agent pack](#)

04

Name: DCF VALUATION

What it does: maps approved forecasts into UFCF, WACC, terminal value, equity bridge, and sensitivities.

When to use: when a defensible forecast and valuation policy exist.

[Download the companion agent pack](#)

05

Name: COMPARABLE COMPANIES

What it does: builds a criteria-led peer set, one aligned data cut, EV bridges, multiples, and outlier notes.

When to use: for a current market view.

[Download the companion agent pack](#)

06

Name: PRECEDENT TRANSACTIONS

What it does: screens transactions, reproduces consideration and multiples, and flags disclosure limits.

When to use: for control-pricing evidence.

[Download the companion agent pack](#)

Use the smallest specialist set that can finish the work.

07-11 of 21 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

07

Name: LBO ANALYSIS

What it does: models sources/uses, debt mechanics, cash sweep, returns, and breakpoints.

When to use: for ability-to-pay support; never to claim financeability.

[Download the companion agent pack](#)

08

Name: THREE-STATEMENT BUILD

What it does: maps approved historicals and drivers into linked statements, supporting schedules, checks, and a change log.

When to use: before downstream DCF, leverage, and scenario analysis.

[Download the companion agent pack](#)

09

Name: SENSITIVITY TABLES

What it does: designs one- and two-variable tests, verifies directionality, and shows which assumptions matter.

When to use: after the base model is stable.

[Download the companion agent pack](#)

10

Name: SCENARIO TOGGLE DESIGN

What it does: links named base/upside/downside assumptions to one control and tests every case without manual overwrites.

When to use: when the team needs alternatives that remain auditable.

[Download the companion agent pack](#)

11

Name: MODEL AUDIT REVIEW

What it does: tests intended use, structure, data, formulas, bridges, circularity, sensitivities, and key recalculations; it never silently fixes.

When to use: after every model build and material revision, before numbers go upstairs.

[Download the companion agent pack](#)

Use the smallest specialist set that can finish the work.

12-16 of 21 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

12

Name: DATA ROOM SYNTHESIS

What it does: inventories approved files into a source ledger, version map, thematic index, and gap list.

When to use: before any diligence analysis begins.

[Download the companion agent pack](#)

13

Name: EARNINGS QUALITY REVIEW

What it does: reconciles reported to adjusted earnings, tests support and recurrence, and flags specialist questions.

When to use: after the source ledger is approved; it does not issue a QoE conclusion.

[Download the companion agent pack](#)

14

Name: RISK FLAG REGISTER

What it does: preserves evidence while mapping impact, likelihood, owner, mitigation, status, and escalation.

When to use: after source-specific findings are reviewed.

[Download the companion agent pack](#)

15

Name: CIM DRAFTING

What it does: creates a document spine, claim map, page briefs, visual specifications, caveats, and HOLD list.

When to use: only after facts and numbers are reviewed.

[Download the companion agent pack](#)

16

Name: MANAGEMENT PRESENTATION

What it does: adapts cleared content into a meeting arc, page briefs, proof map, speaking points, and anticipated Q&A.

When to use: before management rehearsal, inside the approved disclosure scope.

[Download the companion agent pack](#)

Use the smallest specialist set that can finish the work.

17-21 of 21 | Each card maps to a Markdown definition in the companion pack. The playbook explains the route; the pack provides the execution layer.

17

Name: PROCESS LETTER DRAFTING

What it does: assembles approved terms, deadlines, submission mechanics, and counsel placeholders.

When to use: after process owners and counsel define the rules; it never sends.

[Download the companion agent pack](#)

18

Name: BUYER Q&A MANAGEMENT

What it does: assigns stable IDs, deduplicates, routes, drafts from approved evidence, and tracks approval/release status.

When to use: when diligence questions arrive; it never approves or releases.

[Download the companion agent pack](#)

19

Name: DEAL MEMO DRAFTING

What it does: structures evidence, valuation, risks, alternatives, conditions, and unresolved items.

When to use: after model and diligence review; it does not decide.

[Download the companion agent pack](#)

20

Name: COMMITTEE RECOMMENDATION

What it does: presents options, trade-offs, conditions, dependencies, and the exact approval sought.

When to use: after decision owners form a view.

[Download the companion agent pack](#)

21

Name: COMMITTEE Q&A PREP

What it does: builds skeptical questions, evidence-backed draft answers, a weak-point map, and response statuses.

When to use: before committee review; never to invent an answer.

[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