

The Necessary Evil: AI Orchestration for Real Work

Book 4 Companion Resources

Edward Fong

The Necessary Evil Series

necessaryevilbooks.com/go/book4-resources

Version 1.0

Copyright / Use Notice

(c) 2026 Edward Fong. All rights reserved.

These resources are companion worksheets for The Necessary Evil: AI Orchestration for Real Work.

These resources are companion worksheets for The Necessary Evil: AI Orchestration for Real Work. They are planning and review aids, not professional advice. AI can prepare, organize, draft, compare, and check. It cannot own the outcome, approve the work, bear responsibility, or replace qualified human review.

Not legal, financial, medical, security, HR, compliance, or professional advice. Use them with qualified human review where appropriate.

For planning, review, and educational use. No signup required.

Version: Version 1.0

Introduction

This packet contains 16 companion resources for The Necessary Evil: AI Orchestration for Real Work. They help you move from scattered prompting to structured, reviewable, human-governed AI workflows.

Core principle: Human ownership governs. AI specialists assist. Review gates protect the work. Final responsibility remains human.

For updates and additional resources, visit: necessaryevilbooks.com/go/book4-resources

Table of Contents

Section A — Core Framework

- AI Orchestration Quick Framework

- Workflow Setup Checklist

- Human Owner Responsibility Checklist

Section B — Roles and Boundaries

- AI Specialist Role Card Template

- Review Gate Checklist

- Handoff Checklist

Section C — Review and Risk

- Prompt-to-Workflow Conversion Worksheet

- Risk Tier Reference for AI-Orchestrated Work

Section D — Worked Examples

Small Business Orchestration Example

Project Manager Orchestration Example

Systems Administrator / IT Workflow Example

Writer / Content Workflow Example

Section E — Blank Templates

Blank Workflow Canvas

Blank Role Card

Blank Review Gate Sheet

Blank Handoff Checklist

Section A — Core Framework

AI Orchestration Quick Framework

Purpose

A one-page reference for the core roles and structure of a human-governed AI work system. This is the central model taught in Chapter 4 of **The Necessary Evil: AI Orchestration for Real Work**.

When to Use

Use this before starting any multi-step AI-assisted workflow. It reminds you who owns the work, who assists, and where review gates belong. For the full framework explanation, see Chapter 4.

Related resources: Workflow Setup Checklist, Human Owner Responsibility Checklist, Blank Workflow Canvas.

The Framework

Roles

Role	Responsibility
Human Owner	Sets direction, approves output, bears final responsibility
Orchestrator	Coordinates workflow steps, routes tasks, enforces boundaries
Specialist AI	Performs bounded tasks within assigned scope
Reviewer	Checks output for accuracy, tone, completeness, risk
QA / Verification Gate	Confirms work meets standards before handoff

Workflow Structure

Human Owner defines the work

- Orchestrator plans and delegates
 - Specialist AI performs bounded task
 - Reviewer checks output
 - QA gate verifies standards
 - Human Owner approves or returns

Core Principles

1. Human ownership governs every workflow.
2. AI specialists assist within bounded roles.
3. Review gates protect the work before handoff.
4. Final responsibility remains human.
5. No AI role approves, publishes, sends, or commits without human sign-off.

Review Reminder

Before using any AI-prepared output in real work, confirm:

- ☐ A human reviewed the output.
- ☐ The output matches the stated goal.
- ☐ No unsupported claims remain.
- ☐ No sensitive/private data was exposed.
- ☐ The appropriate person approved it.

Workflow Setup Checklist

Purpose

Use before building any AI-assisted workflow. Ensures you have defined the goal, owner, boundaries, roles, and review gates before AI begins work.

When to Use

At the start of any new multi-step AI workflow, project, or task delegation.

Related resources: AI Orchestration Quick Framework, AI Specialist Role Card Template, Blank Workflow Canvas.

Checklist

- ☐ **Goal defined:** What is the final deliverable?
- ☐ **Owner identified:** Who bears final responsibility?
- ☐ **Source material identified:** What inputs does the AI need?
- ☐ **Constraints stated:** What must NOT be done?
- ☐ **Scope boundaries set:** Where does AI authority stop?
- ☐ **Specialist roles assigned:** Which AI roles are needed?
- ☐ **Review gates placed:** Where must a human check output?
- ☐ **Final output defined:** What format/artifact is expected?
- ☐ **Handoff action defined:** What happens after the workflow completes?
- ☐ **Escalation path defined:** What happens if something goes wrong?

How to Use

1. Fill in each checkbox before delegating work to AI.
2. If any item is blank, the workflow is not ready to start.
3. Share the completed checklist with anyone who will review the AI output.

Review Reminder

A workflow without a defined owner, review gate, or escalation path is not governed. Do not start.

Human Owner Responsibility Checklist

Purpose

Clarifies what must remain under human ownership in any AI-assisted workflow. Use to prevent accidental delegation of judgment, ethics, or accountability to AI.

When to Use

Before finalizing any workflow. Review whenever AI output will be published, sent, submitted, or used in a decision.

Related resources: Review Gate Checklist, Risk Tier Reference.

What Humans Must NOT Delegate to AI

- ☐ Final judgment on whether output is acceptable
- ☐ Ethical decisions about fairness, harm, or impact
- ☐ Truthfulness verification for claims, facts, and data
- ☐ Stakeholder impact assessment (who is affected and how)
- ☐ Publishing / sending decisions (when and to whom)
- ☐ Legal decisions (contracts, compliance, liability)
- ☐ Financial decisions (pricing, budget, payment)
- ☐ Medical / health decisions
- ☐ Security / access decisions
- ☐ HR / personnel decisions
- ☐ Accountability for consequences of the work

What AI May Assist With

- Drafting
- Organizing
- Comparing options
- Summarizing
- Checking consistency

- Formatting
- Research support

Key Principle

AI prepares. Humans decide.

If you cannot explain why the output is acceptable, you have not reviewed it.

Section B — Roles and Boundaries

AI Specialist Role Card Template

Purpose

Define a bounded AI specialist role before assigning it work. Prevents scope creep, unauthorized actions, and unclear accountability.

When to Use

Whenever you create a new AI role in a workflow. Fill one card per specialist. For a blank fillable version, see the Blank Role Card template.

Related resources: AI Orchestration Quick Framework, Blank Role Card.

Role Card

Field	Your Entry
Role name	
Task boundary	What this role does (and nothing else)
Allowed inputs	What information/context it receives
Prohibited actions	What it must NOT do
Expected output	What artifact it produces
Output format	Text / table / checklist / draft / other
Review requirement	Who reviews this role's output
Escalation trigger	When must this role stop and ask a human
Quality standard	How output is judged acceptable

Example

Field	Entry
Role name	Research Summarizer
Task boundary	Summarize provided sources into key findings
Allowed inputs	Source documents provided by owner

Prohibited actions	Must not add claims not in sources; must not recommend actions
Expected output	Bullet-point summary with source attribution
Output format	Markdown list
Review requirement	Human Owner reviews before use
Escalation trigger	If sources conflict or contain sensitive claims
Quality standard	Every bullet traceable to a source

How to Use

1. Fill one card per AI specialist role in your workflow.
2. Share the card with anyone who will use or review the role's output.
3. If the role produces output outside its boundary, the workflow has a governance gap.

Review Gate Checklist

Purpose

Use at every review gate in an AI-assisted workflow. Ensures output is checked for accuracy, tone, risk, and completeness before it advances.

When to Use

After any AI specialist produces output and before that output moves to the next step, is published, or is sent to a stakeholder. For a blank fillable version, see the Blank Review Gate Sheet.

Related resources: Human Owner Responsibility Checklist, Risk Tier Reference, Blank Review Gate Sheet.

Checklist

- ☐ **Accuracy:** Are all facts, numbers, and claims correct?
- ☐ **Tone:** Does the output match the intended voice and audience?
- ☐ **Completeness:** Is anything missing that was requested?
- ☐ **Missing context:** Does the output assume information the reader won't have?
- ☐ **Unsupported claims:** Are there statements without evidence or source?
- ☐ **Risk:** Could this output cause harm, confusion, or liability if used as-is?
- ☐ **Audience fit:** Is this appropriate for the intended recipient?
- ☐ **Sensitive content:** Does it expose private, confidential, or personal information?
- ☐ **Actionability:** Can the recipient use this output without further clarification?
- ☐ **Qualified review trigger:** Does this topic require a specialist human reviewer (legal, medical, financial, security)?

Decision After Review

Outcome	Action
All checks pass	Advance to next step
Minor issues found	Return to AI with correction instructions
Major issues found	Escalate to Human Owner
Qualified review needed	Route to appropriate specialist human

Unsafe or harmful	Stop workflow immediately
-------------------	---------------------------

Handoff Checklist

Purpose

Use when transferring AI-prepared work to a human reviewer, stakeholder, or next workflow step. Ensures nothing is sent prematurely or without appropriate context.

When to Use

At any handoff point: AI-to-human, step-to-step, draft-to-review, or internal-to-external. For a blank fillable version, see the Blank Handoff Checklist.

Related resources: Review Gate Checklist, Blank Handoff Checklist.

Checklist

- ☐ **What is ready:** Clearly state what the deliverable is
- ☐ **What remains uncertain:** Flag anything unverified or assumed
- ☐ **Who must review:** Name the responsible reviewer
- ☐ **What should NOT be sent/published yet:** Identify anything premature
- ☐ **Context provided:** Does the recipient have enough background?
- ☐ **Format correct:** Is the output in the expected format?
- ☐ **Sensitive content flagged:** Is anything confidential, personal, or risky?
- ☐ **Final responsibility owner identified:** Who is accountable for what happens next?
- ☐ **Escalation path clear:** What happens if the recipient finds a problem?
- ☐ **Approval recorded:** Has the required approval been given?

Handoff Statement Template

Deliverable: [what is being handed off]

Status: [draft / reviewed / approved]

Reviewed by: [name or role]

Remaining concerns: [list or "none"]

Next action needed: [what the recipient should do]

Do NOT: [what should not happen yet]

Final responsibility: [name or role]

Section C — Review and Risk

Prompt-to-Workflow Conversion Worksheet

Purpose

Turn a single AI prompt into a structured, reviewable workflow with defined steps, roles, and human approval.

When to Use

When you catch yourself writing a long, complex prompt and realize the task needs structure, not just a better prompt. This worksheet bridges Book 2's prompting framework and Book 4's orchestration model.

Related resources: AI Orchestration Quick Framework, Workflow Setup Checklist.

Worksheet

Step 1 — Define the Goal

What is the final deliverable?

Goal: _____

Step 2 — Identify Inputs

What information does the AI need to start?

Input 1: _____

Input 2: _____

Input 3: _____

Step 3 — Break Into Workflow Steps

Convert your prompt into 3 bounded steps:

Step	Task	AI Role	Output
1			
2			
3			

Step 4 — Assign One Specialist Role

Companion worksheet — not professional advice. Final responsibility remains human. Version 1.0 — necessaryevilbooks.com/go/book4-resources —

Pick the most important step and define a role card:

Role name: _____

Task boundary: _____

Prohibited actions: _____

Expected output: _____

Step 5 — Place One Review Gate

Where must a human check before the work advances?

Review gate after step: _____

Reviewer: _____

What they check: _____

Step 6 — Define Final Human Action

What does the human do after the workflow completes?

Final action: _____

Approved by: _____

Result

You now have a 3-step workflow with a defined role, review gate, and human approval instead of a single unstructured prompt.

Risk Tier Reference for AI-Orchestrated Work

Purpose

Classify AI-assisted tasks by risk level. Higher-risk tasks require more review, more human involvement, and stricter boundaries.

When to Use

Before assigning work to AI. Check the risk tier to determine how much human oversight is needed. See Chapters 12–13 for the full verification and governance model.

Related resources: Review Gate Checklist, Human Owner Responsibility Checklist.

Risk Tiers

Tier 1 — Low Risk / Internal Organization

Characteristic	Example
Impact if wrong	Minor inconvenience
Audience	Yourself only
Stakes	Low
Review needed	Light self-review

Examples: personal notes, brainstorming lists, internal drafts, file organization.

Tier 2 — Moderate Risk / Drafting for Review

Characteristic	Example
Impact if wrong	Requires revision
Audience	Internal team or trusted reviewer
Stakes	Moderate
Review needed	Peer review before use

Examples: draft emails, internal reports, meeting summaries, project plans.

Tier 3 — High Risk / Stakeholder Impact

Characteristic	Example
----------------	---------

Impact if wrong	Damages reputation, relationship, or trust
Audience	External stakeholders, clients, public
Stakes	High
Review needed	Formal review gate + domain expert

Examples: client communications, published content, business proposals, public statements.

Tier 4 — Qualified Review Required

Characteristic	Example
Impact if wrong	Legal, financial, medical, or safety harm
Audience	Regulated or vulnerable populations
Stakes	Very high
Review needed	Qualified professional review mandatory

Examples: legal documents, financial advice, medical information, security configurations, compliance materials.

Tier 5 — Human Only / Do Not Delegate

Characteristic	Example
Impact if wrong	Irreversible harm
Audience	Anyone
Stakes	Critical
Review needed	Human judgment only — AI must not perform

Examples: final legal approval, medical treatment decisions, emergency response, personnel termination, security incident response execution.

How to Use

1. Identify the risk tier for your task.
2. Match the review intensity to the tier.
3. Never use Tier 1 review for Tier 3+ work.

4. When in doubt, escalate to the next higher tier.

Section D — Worked Examples

Small Business Orchestration Example

Purpose

Shows how a small-business owner can use a simple AI-orchestrated workflow to prepare a customer proposal.

Scenario

A freelance consultant needs to prepare a project proposal for a potential client. This example shows how even a simple 5-step workflow keeps the human in control of what gets sent.

Workflow

Step	Role	Task	Output
1	Human Owner	Define project scope, budget range, client context	Scope brief
2	Research Specialist	Summarize client's industry and comparable project pricing	Research summary
3	Draft Specialist	Write proposal draft using scope brief + research	Proposal draft
4	Reviewer	Check for unsupported claims, pricing accuracy, tone	Review notes
5	Human Owner	Revise, approve, and send to client	Final proposal

Role Card: Draft Specialist

Field	Value
Role name	Proposal Draft Specialist
Task boundary	Draft proposal from provided scope and research
Prohibited actions	Must not invent pricing; must not promise deliverables not in scope

Expected output	1–2 page proposal draft
Review requirement	Human Owner reviews before sending

Review Gate

After Step 3, Human Owner checks:

- Does the proposal match the agreed scope?
- Are pricing claims accurate?
- Is the tone appropriate for the client?
- Are there promises the consultant cannot keep?

Key Principle

The AI drafts. The owner decides what to send.

Project Manager Orchestration Example

Purpose

Shows how a project manager can use AI-orchestrated workflow to prepare a weekly status report.

Scenario

A PM manages a 5-person team and needs to prepare a weekly status report for leadership. This example shows how AI can organize inputs without inventing progress or hiding problems.

Workflow

Step	Role	Task	Output
1	Human Owner (PM)	Collect team updates via standup notes	Raw status inputs
2	Summary Specialist	Organize updates by workstream, flag blockers	Structured summary
3	Risk Specialist	Identify risks and suggest escalation items	Risk assessment
4	Draft Specialist	Compose status report for leadership audience	Report draft
5	Reviewer	Check accuracy against raw inputs, verify tone	Review notes
6	Human Owner (PM)	Approve and send to leadership	Final report

Role Card: Summary Specialist

Field	Value
Role name	Status Summary Specialist
Task boundary	Organize raw team updates into workstream categories
Prohibited actions	Must not invent progress not reported; must not hide blockers
Expected output	Categorized summary with blockers highlighted
Review requirement	PM reviews before report drafting

Review Gate

After Step 4, PM checks:

- Does the report accurately reflect team status?
- Are blockers correctly identified?
- Is the tone appropriate for leadership?
- Are there claims not supported by team input?

Key Principle

AI organizes and drafts. The PM owns the accuracy and sends the report.

Systems Administrator / IT Workflow Example

Purpose

Shows how an IT administrator can use AI-orchestrated workflow to prepare a change request package for review. Emphasizes that AI prepares documentation — it does not execute changes.

Scenario

A sysadmin needs to prepare a change request for a server patch deployment.

Workflow

Step	Role	Task	Output
1	Human Owner (Sysadmin)	Define change scope, affected systems, maintenance window	Change brief
2	Documentation Specialist	Draft change request using organizational template	Change request draft
3	Rollback Planning Specialist	Draft rollback plan based on change scope	Rollback plan draft
4	Reviewer	Verify change request completeness, rollback feasibility	Review notes
5	Human Owner (Sysadmin)	Submit change request through organizational approval process	Submitted CR

Role Card: Documentation Specialist

Field	Value
Role name	Change Request Documentation Specialist
Task boundary	Draft change request document from provided change brief
Prohibited actions	Must NOT execute commands; must NOT access systems; must NOT approve changes; must NOT bypass change control
Expected output	Formatted change request draft

Review requirement	Sysadmin reviews; change advisory board approves
Escalation trigger	If change affects production systems without tested rollback

Review Gate

After Step 3, Sysadmin checks:

- Is the change scope accurate?
- Is the rollback plan feasible and tested?
- Are affected systems correctly identified?
- Does the change follow organizational policy?
- Has the maintenance window been approved?

Critical Boundaries

- AI prepares documentation only. AI does NOT execute changes.
- Change control, authorization, and execution remain human responsibilities.
- Rollback plans must be validated by a qualified human before use.
- Security patches require organizational approval regardless of AI assistance.
- No credentials, access tokens, or system passwords in AI context.

Key Principle

AI drafts the paperwork. Humans authorize, execute, and verify the change.

Writer / Content Workflow Example

Purpose

Shows how a writer or content creator can use AI-orchestrated workflow to produce a reviewed article draft.

Scenario

A freelance writer needs to produce a 1,500-word article on a topic assigned by an editor. This example shows how AI can assist with research and structure while the writer retains voice, accuracy, and originality.

Workflow

Step	Role	Task	Output
1	Human Owner (Writer)	Define topic, angle, audience, word count, sources	Creative brief
2	Research Specialist	Gather key points from provided sources	Research notes
3	Outline Specialist	Create article outline from brief + research	Outline
4	Draft Specialist	Write full article draft from outline	Article draft
5	Reviewer	Check facts, tone, word count, unsupported claims	Review notes
6	Human Owner (Writer)	Revise in own voice, verify accuracy, submit to editor	Final article

Role Card: Draft Specialist

Field	Value
Role name	Article Draft Specialist
Task boundary	Write article from approved outline and research notes
Prohibited actions	Must not fabricate quotes; must not add unsourced claims; must not plagiarize
Expected output	1,500-word article draft

Review requirement	Writer reviews for voice, accuracy, and originality
--------------------	---

Review Gate

After Step 4, Writer checks:

- Does the draft match the assigned angle?
- Are all claims supported by the provided sources?
- Is the tone appropriate for the target publication?
- Does it sound like the writer's voice (not generic AI)?
- Are there any fabricated quotes or unsourced statistics?

Key Principle

AI assists with research and structure. The writer owns the voice, accuracy, and final submission.

Originality note: AI-drafted text should always be checked for unintentional plagiarism or overly generic phrasing. The writer's job includes making the final output distinctly theirs.

Section E — Blank Templates

Blank Workflow Canvas

Purpose

A blank template for designing any AI-orchestrated workflow from scratch.

Instructions

Fill in each row with one step of your workflow. Add or remove rows as needed. See the AI Orchestration Quick Framework for role definitions.

Related resources: AI Orchestration Quick Framework, Workflow Setup Checklist, AI Specialist Role Card Template.

Workflow Canvas

Step	Task Description	AI Role	Input	Output	Review Gate?
1					Yes / No
2					Yes / No
3					Yes / No
4					Yes / No
5					Yes / No
6					Yes / No

Workflow Metadata

Field	Your Entry
Workflow name	
Goal	
Human Owner	
Final output	
Final approval by	

Blank Role Card

Purpose

A blank template for defining any AI specialist role.

Instructions

Fill one card per AI role in your workflow.

Role Card

Field	Your Entry
Role name	
Task boundary	
Allowed inputs	
Prohibited actions	
Expected output	
Output format	
Review requirement	
Escalation trigger	
Quality standard	

Blank Review Gate Sheet

Purpose

A blank template for documenting a review gate decision.

Instructions

Fill after reviewing AI output at any gate point.

Review Gate Record

Field	Your Entry
Workflow name	
Step reviewed	
Reviewer	
Date	

Checks

- ☐ Accuracy verified
- ☐ Tone appropriate
- ☐ Completeness confirmed
- ☐ No unsupported claims
- ☐ No missing context
- ☐ Risk assessed
- ☐ Audience fit confirmed
- ☐ No sensitive data exposed
- ☐ Qualified review not needed / triggered

Decision

Decision	✓
Advance to next step	
Return for revision	
Escalate to Human Owner	
Route to qualified reviewer	
Stop workflow	

Notes

Blank Handoff Checklist

Purpose

A blank template for documenting any handoff point.

Instructions

Fill when transferring work from one step/person to the next.

Handoff Record

Field	Your Entry
Deliverable	
Status (draft / reviewed / approved)	
Prepared by	
Handed off to	
Date	

Checks

- ☐ What is ready is clearly stated
- ☐ What remains uncertain is flagged
- ☐ Reviewer identified
- ☐ Nothing premature is being sent
- ☐ Recipient has enough context
- ☐ Format is correct
- ☐ Sensitive content flagged if present
- ☐ Final responsibility owner identified
- ☐ Escalation path clear

☐ Approval recorded if required

Do NOT

Next Action for Recipient
