

Multi-Model Collaboration Board Starter Kit

A no-code starting point for coordinating AI tools in one shared workspace

Bryon Stafford

The operating idea

I now run almost everything important through both models because the finished work is noticeably better.

For each task, assign one AI tool to create and a different tool to check the work independently. The roles can switch on the next task. A shared workspace preserves context, while the collaboration board records ownership, handoffs, disagreements, and human decisions.

If the models cannot agree, the human owner decides.

Start with one task, one board, and one clean handoff

You do not need a dashboard, automation, or new software on day one. Choose one real task you already understand and create one shared file called `CURRENT-STATE.md`.

Choose the participating AI tools, then assign the Creator and Independent Checker roles for this task. Roles can switch on the next task. If you do not know which tool should do which job, ask before starting. Think of `CURRENT-STATE.md` as the traffic director: it tells each tool whose turn it is, what happened previously, what it should do next, and when to hand the work back. The board coordinates the tools; it does not open or wake separate AI applications.

The minimum board needs to answer:

- What is active?
- Who is creating and who is checking?
- What is its status?
- What happens next?

COPY-AND-PASTE PROMPT: CREATE THE ONE-FILE BOARD

Create one simple collaboration board in my shared workspace.

Before changing any files, ask which AI tools will participate and who will be the Creator and Independent Checker for the first task. If I do not know, make a recommendation and wait for my confirmation.

Then show me the proposed file location and wait for approval.

Create:
Agent-Collaboration/CURRENT-STATE.md

Use three sections:

1. Active claims
2. Waiting on me
3. Recently closed

Use this Active claims table:

Work item | Creator | Independent Checker | Status | Next action | Updated

Keep the file short. State that roles are assigned per task and may switch on the next task. Do not add a dashboard, automation, integration, archive, or extra protocol file yet.

BLANK CURRENT-STATE.MD TEMPLATE

```
# Collaboration Board - Current State

*Read before acting. Claim before editing.*
*Creator and Independent Checker are assigned per task and may switch next time.*
*Updated: YYYY-MM-DD HH:MM - Agent*

## Active claims

| Work item | Creator | Independent Checker | Status | Next action | Updated |
|---|---|---|---|---|---|
| Example client brief - `brief.md` | Codex | Claude | Creator drafting | Codex finishes and releases; human relays to Claude | YYYY-MM-DD |

## Waiting on me

| Item | What I need to decide | Agent follow-through | Updated |
|---|---|---|---|
| Example visual choice | Select A or B | Assigned model issues the final version | YYYY-MM-DD |

## Recently closed

- None yet.
```

What one turn looks like

Every working turn follows the same small, auditable cycle.

- 1. Read the board.** Recover the current state.
- 2. Check overlap.** See who is already working.
- 3. Claim a lane.** Name the task, assigned role, and exact file scope.
- 4. Do the work.** Stay inside the claim.
- 5. Update the row.** Record the result, uncertainty, and next action.
- 6. Hand it back.** Release the lane cleanly.

The handoff is written for whoever comes next: you, the same model tomorrow, or a different model five minutes later. With separate AI applications, the human owner carries the handoff when one tool cannot wake the other.

Five rules keep the system reliable

01 - START WITH THE BOARD

Read the current state before doing anything. Do not rely on chat memory to recover ownership or status.

02 - CLAIM EXACTLY WHAT YOU OWN

Name the task, file scope, and owner. Concurrent work is allowed only in clearly non-overlapping files.

03 - DELEGATE INTELLIGENTLY

Use a strong model to manage and lower-cost subagents for large, independent work when that division is worthwhile. Delegation does not automatically reduce total tokens because every worker needs startup context. Small jobs stay with one model. Shared-file edits always stay one at a time.

04 - LEAVE A COMPLETE HANDOFF

Record what changed, what remains uncertain, who acts next, and what that person or model should do.

05 - THE HUMAN DECIDES

If the models cannot agree, the human owner decides. Material uncertainty and irreversible actions also stop for human review.

Better work through a challenge loop

The second model is not there to politely agree. Its job is to look for errors, omissions, contradictions, weak evidence, and unsafe assumptions.

```
CREATE      First model drafts and shows its evidence
|
CHALLENGE   Independent model hunts for flaws
|
COMPARE     Every meaningful difference is logged
|
RECONCILE   Both models return to the source evidence
|
RECHECK     The disputed point is tested again
|
+----- Loop through compare, reconcile, and recheck when needed
```

The loop ends in one of three states:

- **AGREE** - the evidence supports a shared answer.
- **ELEVATE** - the models cannot resolve a judgment call, so the human decides.
- **RFI** - the available documents cannot support a safe answer, so the question remains open.

Agreement alone is not proof. The evidence must support the result.

COPY-AND-PASTE CHECKER PROMPT

```
Review this completed work as an independent checker.

Your job is to find what is wrong with it. Check the same source material and list
errors, omissions, contradictions, weak evidence, and unsafe assumptions. Do not
summarize the work and do not praise it.

For every disagreement, show:
- the disputed claim or result;
- the supporting source or evidence;
- your recommended correction;
- whether the outcome is AGREE, ELEVATE, or RFI.

If the evidence does not settle the difference, say so plainly.
```

How the board and dashboard work

The collaboration board is the source of truth that the AIs read and update. A dashboard is an optional, read-only, human-facing status window that helps the owner see where the process stands without reading the entire board.

Part	Primary job	Source of truth?	Primary user
Shared workspace	Stores instructions, sources, decisions, and deliverables	Yes	People and AI tools
CURRENT-STATE.md	Records ownership, status, next action, and decisions	Yes	People and AI tools
Workspace dashboard	Makes claims, decisions, and routines easy to scan	No - generated snapshot	Human owner

Add a dashboard only after the manual board is trusted.

OPTIONAL DASHBOARD PROMPT

Using CURRENT-STATE.md as the canonical source, create a local, read-only HTML dashboard that serves as a human-facing status window and makes ownership, next actions, human decisions, and recurring routines easy to scan.

Requirements:

- The dashboard is a generated snapshot, never the source of truth.
- Show the last-updated time for every source.
- Warn when a source is stale.
- Keep private data local.
- Include a presentation-safe mode using only allowlisted rows.
- Do not add editing controls or external hosting.
- Use another AI model to check parsing, dates, links, privacy, and failure behavior.

Build your personal collaboration board

1. **Choose one real task.** Start with work you already understand.
2. **Create one board.** `CURRENT-STATE.md` is enough.
3. **Choose Creator and Checker.** Use the best-fit model as Creator and a different model as Independent Checker. Roles can switch next time; if you are unsure, ask.
4. **Define the handoff.** Decide what gets claimed, updated, and released.
5. **Test one real job.** Run one complete handoff from claim through release.
6. **Add a dashboard later.** Only after the board is trusted.

Make it yours

Start with the plain version. Change only the parts that solve a problem you can name.

WHAT TO CUSTOMIZE FIRST

- **Your AI tools.** Claude, Codex, Gemini, Grok, or any combination that works for you. Assign roles per task instead of permanently tying one tool to creating or checking.
- **Your columns.** Start with Creator, Independent Checker, status, and next action; add scope or risk only when needed.
- **Your checker rule.** Example: anything a client sees gets an independent second look.
- **Your escalation line.** Decide when the AIs stop and bring the question to you.

PRACTICAL USES

- **Newsletters and content:** one drafts; another checks names, dates, and links.
- **Client deliverables:** an independent second read before anything goes out.
- **Research:** one gathers sources; another challenges the claims and evidence.
- **Stop-and-start projects:** the board makes restarting clean and fast.
- **Extraction work:** the models compare results, return to the evidence, and loop until they agree or elevate the question.
- **Important decisions:** one model develops the recommendation; another pressure-tests its assumptions, evidence, and risks.
- **Security and workflow audits:** one model checks the other model's cleanup, controls, and process changes for anything missed.

Grow the protocol only when needed

Once the one-file board is working reliably, add only the component that solves the next operating problem.

```
Agent-Collaboration/  
  CURRENT-STATE.md      # live board and human decisions  
  README.md             # detailed operating rules  
  COLLABORATION-LOG.md  # long reasoning and decision history  
  BOARD-ARCHIVE.md      # exact history of closed work
```

ADVANCED BOARD COLUMNS

When ownership becomes more complex, expand Active claims to:

```
Work item | Creator | Independent Checker | Status | File/area scope | Structural? | Next  
action | Updated
```

ADVANCED PROTOCOL RULES

- Create, delete, rename, move, same-file edit, and overwrite operations are exclusive.
- Parallel work is allowed only on clearly non-overlapping files.
- A stale claim never releases itself; reconcile it with the owner or human conductor.
- Put the next actor and next action in every row.
- Replace the complete row and verify its column count after an update.
- Put long reasoning in the collaboration log, not the live board.
- Archive closed work without deleting its history.
- Keep irreversible actions under human control.

Safety checklist

- The board remains authoritative if the dashboard disagrees.
- Never expose credentials or secrets.
- Do not publish a live internal board without reviewing every row.
- Use an allowlisted presentation mode for demonstrations.
- Keep client, employee, student, health, legal, and personal information out of public screenshots unless specifically cleared.
- Preserve earlier versions; do not silently overwrite them.
- Test failure behavior, not only the happy path.

Useful upgrades after the basics work

- Stale-claim warnings
- A compact Waiting on Me view
- A daily what-changed digest
- Routine-health alerts
- Per-domain filtered views generated from the same board
- Handoff-cycle and review-latency reporting after statuses become structured

Start with a real task. Keep the board short. Add more only after something real becomes hard to manage.

Appendix - Copy-and-Paste Toolkit

How to use this toolkit

Copy one prompt at a time into the AI tool that should perform that role. Replace every bracketed placeholder before sending it.

Creator and Independent Checker are task-specific roles, not permanent model identities. Claude might create and Codex might check one task; the roles may reverse on the next task.

If you do not know which tool should do which job, start with Prompt 1 and ask. The board coordinates the tools, but it does not open, launch, or wake separate AI applications. When the tools cannot contact each other, the human owner carries the handoff.

Think of CURRENT-STATE.md as the traffic director: it tells each tool whose turn it is, what happened previously, what it should do next, and when to hand the work back.

Task assignment card

Complete this card before starting the task. Put the same assignment in the collaboration board.

```
Task:
Participating AI tools:
Creator for this task:
Independent Checker for this task:
Human decision-maker:
Shared workspace or source location:
Files the Creator may change:
Files the Checker may change:
```

Prompt 1 - Recommend the Creator and Checker

Use this when you know the task and available tools but do not know how to assign the roles.

```
I need to complete this task:

[DESCRIBE THE TASK]

The AI tools available to me are:

[LIST THE TOOLS, SUCH AS CLAUDE, CODEX, CHATGPT, GEMINI, OR GROK]

Recommend which tool should be the Creator and which should be the Independent
Checker for this task. Briefly explain the assignment based on the tools'
strengths, access to the required information, and ability to work with the
needed files.

Creator and Checker are task-specific roles. Do not assume that a particular
tool always holds the same role. If two assignments are equally reasonable,
show both and explain the tradeoff.

Do not begin the task until I confirm the roles.
```

Prompt 2 - Create the one-file collaboration board

Use this once to create the minimum working board. If the board already exists, do not recreate it.

Create one simple collaboration board in my shared workspace.

Before changing any files, ask which AI tools will participate and who will be the Creator and Independent Checker for the first task. If I do not know, make a recommendation and wait for my confirmation.

Then show me the proposed file location and wait for approval.

Create:

Agent-Collaboration/CURRENT-STATE.md

Use three sections:

1. Active claims
2. Waiting on me
3. Recently closed

Use this Active claims table:

Work item	Creator	Independent Checker	Status	Next action	Updated
-----------	---------	---------------------	--------	-------------	---------

Keep the file short. State that roles are assigned per task and may switch on the next task. Do not add a dashboard, automation, integration, archive, or extra protocol file yet.

Complete `CURRENT-STATE.md` template

Save the contents of this code block as `Agent-Collaboration/CURRENT-STATE.md`.

```
# Collaboration Board - Current State

*Read before acting. Claim before editing.*
*Creator and Independent Checker are assigned per task and may switch next time.*
*Updated: YYYY-MM-DD HH:MM - Agent*

## Active claims

| Work item | Creator | Independent Checker | Status | Next action | Updated |
|---|---|---|---|---|---|
| Example client brief - `brief.md` | Codex | Claude | Creator drafting | Codex finishes and releases; human relays to Claude | YYYY-MM-DD |

## Waiting on me

| Item | What I need to decide | Agent follow-through | Updated |
|---|---|---|---|
| Example visual choice | Select A or B | Assigned model issues the final version | YYYY-MM-DD |

## Recently closed

- None yet.
```

Prompt 3 - Activate the Creator

Send this to the AI tool assigned as Creator for the current task.

```
You are the Creator for this task.

Task: [TASK]
Creator: [CREATOR TOOL]
Independent Checker: [CHECKER TOOL]
Human decision-maker: [NAME]
Shared board: Agent-Collaboration/CURRENT-STATE.md

Before working:
1. Read the collaboration board.
2. Confirm that you are assigned as Creator for this task.
3. Check for overlapping work.
4. Claim the Creator lane and identify the exact files you may change.

Complete the work inside that scope. Show the source or evidence supporting
important claims. Record uncertainty instead of hiding it.

When the draft is ready:
- stop changing it;
- update the board status to READY FOR CHECKER;
- record what changed, the evidence used, known uncertainty, and the exact next
  action for the Independent Checker;
- tell the human owner what message to relay to the Checker.

Do not perform the independent check on your own work.
```

Prompt 4 - Activate the Independent Checker

Send this to the different AI tool assigned as Independent Checker.

```
You are the Independent Checker for this task.

Task: [TASK]
Creator: [CREATOR TOOL]
Independent Checker: [CHECKER TOOL]
Human decision-maker: [NAME]
Shared board: Agent-Collaboration/CURRENT-STATE.md

Before reviewing:
1. Read the collaboration board.
2. Confirm that the Creator released the work for review.
3. Claim the Checker lane and identify the files you may change.
4. Read the original source material, not only the Creator's summary.

Your job is to find what is wrong with the work. Check for errors, omissions,
contradictions, weak evidence, and unsafe assumptions. Do not merely agree with
the Creator and do not rewrite the work unless that is explicitly assigned.

For every meaningful disagreement, show:
- the disputed claim or result;
- the supporting source or evidence;
- your recommended correction;
- whether the current outcome is AGREE, ELEVATE, or RFI.

Update the board with the findings, remaining uncertainty, and exact next
action. If the Creator must respond, tell the human owner what message to relay.
```

Prompt 5 - Reconcile and recheck a disagreement

Use this after the Checker records one or more disagreements.

Read Agent-Collaboration/CURRENT-STATE.md and proceed from the recorded next action for this task:

[TASK]

Read the original source material and stay in your currently assigned role.

If you are the Creator:

- answer every Checker disagreement separately;
- return to the source evidence;
- accept supported corrections;
- explain any disagreement that remains;
- update the work and board, then release it for recheck.

If you are the Independent Checker:

- recheck every Creator response against the source;
- confirm which corrections are resolved;
- identify what remains unsupported;
- update the board with AGREE, ELEVATE, or RFI.

Do not settle a difference by averaging, voting, or splitting the difference. Loop through response and recheck when new evidence can resolve the issue. If the evidence cannot resolve a judgment call, elevate it to the human owner.

Prompt 6 - Human relay between separate AI applications

Use the appropriate message when one AI application cannot contact or wake the other.

Send to the Checker when the Creator finishes:

Please read Agent-Collaboration/CURRENT-STATE.md and claim the Independent Checker lane for [TASK]. Review the Creator's released work and follow the next action recorded on the board.

Send to the Creator when the Checker finds issues:

Please read Agent-Collaboration/CURRENT-STATE.md and proceed from the recorded next action for [TASK]. Claim the Creator lane, respond to every Checker finding, and follow the reconciliation steps recorded on the board.

Use when the board is waiting on the human:

Please read the Waiting on me section of Agent-Collaboration/CURRENT-STATE.md. Summarize only the decisions that require me, explain the consequence of each choice, and wait for my decision before continuing.

Prompt 7 - Create the optional dashboard

Use this only after the manual board is working reliably.

Using CURRENT-STATE.md as the canonical source, create a local, read-only HTML dashboard that serves as a human-facing status window and makes ownership, next actions, human decisions, and recurring routines easy to scan.

Requirements:

- The dashboard is a generated snapshot, never the source of truth.
- Show the last-updated time for every source.
- Warn when a source is stale.
- Keep private data local.
- Include a presentation-safe mode using only allowlisted rows.
- Do not add editing controls or external hosting.
- Use another AI model to check parsing, dates, links, privacy, and failure behavior.

Quick operating reminder

Assign the best-fit tool to create and a different tool to check the work. Roles may switch on the next task. If you do not know which tool should do which job, ask before starting.