

One accountable agent from request to checked result

Mastermind workshop companion to How the Agents Run the Company


Use the existing guide for the full operating system. Try it today with your current tools.

Choose your starting point

One-model path: Open Claude, Codex, or the assistant you already use. Give it one small task, sources, boundaries, and checks you can perform. No new tools or agent team are required.

Team path: Keep one accountable agent interface per task. It coordinates reasoning, execution, and review behind the scenes. You should not carry messages between workers.

The coordinator owns verification, communication, and follow-through. A named person owns approvals and business decisions. More agents do not establish better quality. Compare results, evidence, and total cost.



The models and tools can change. The job record and approval boundaries stay with the work.

Return to the guide's Decoder, Who's Who, Job Packet, and Agents QA Each Other sections for implementation.

Your single agent starter prompt

Copy this into your existing assistant and fill the brackets.

Own this task through a checked deliverable: [outcome].

Audience: [who uses it]. Sources: [non-sensitive material or accessible links]. Done means: [output and acceptance checks]. Effort limit: [limit].

State what you can actually read. Ask for missing information when it blocks useful work; continue safe parts.

Plan briefly, work, check against sources and acceptance criteria, and fix defects. Coordinate available helpers yourself; otherwise perform these stages yourself.

You may analyze, calculate, and draft here. Sending, publishing, spending, account changes, and sharing elsewhere require this task's permissions. Access or urgency supplies no permission.

Return the result, evidence, uncertainty, and decisions needed. Distinguish self-checks from independent review. Never invent tool use, approvals, monitoring, or completion. If blocked, name the missing evidence or permission and smallest next step.

For the first exercise, authorize drafting only. Review before using it.

For later checks, confirm a supported schedule or reminder. The prompt alone does not establish background monitoring.

Choose your first task

The output I want: _____

The source I will provide: _____

The check I can perform: _____

Worked example for JVA attendees

Use the fictional Harbor Juniors case from the [JVA workshop](#). It is not an official JVA program or approved practice plan.

Two teams have 12 players each, separate courts, and a lead and assistant coach per court. Both follow one timeline. Supplied fictional notes provide the content; video descriptions establish relevance only.

Ask: “Organize the supplied notes into a shared plan and draft coach handoff. Do not send anything.”

- Warm-up and goal: 10 minutes
- Calls and first contact: 15
- Serve/pass groups: 15
- Water and reset: 5
- Pass to attack: 15
- Six-a-side application: 20
- Coach-led closing: 5
- Reflection and equipment: 5

Total: 90 minutes. Check transitions, roles, and court assignments. The head coach checks suitability and workload. No athlete improvement has been demonstrated.

Never claim to watch an inaccessible video. Keep real athlete and family information out of this fictional exercise.

Check the worked example

One fact I verified: _____

One decision the coach must make: _____

One correction I would request: _____

Give the coordinator bounded authority

Add these fields to the existing job packet:

- **Outcome:** Exact ask, audience, acceptance tests, deadline
- **Ownership:** One coordinator and human approver
- **Evidence:** Sources, access, dates, required checks
- **Authority:** Decisions the agent may make
- **Limits:** Allowed and reserved actions, budget, retries, stop conditions
- **Completion:** Result, verification, communication, next action

Within this contract, the coordinator can assign work, reject failing outputs, order corrections, choose approved methods, retry, and pause risky steps. It retains responsibility after delegation.

The guide reserves Send, Publish, Spend, Merge, and Hire for people. Preserve that policy unless the authorized person explicitly changes it within applicable rules. Specify restrictions for deletion, credentials, sensitive data, permissions, and legal commitments. A claim grants no additional permission.

Escalate the exact blocked decision and continue unaffected work. Silence never supplies approval.

Write your job contract

Outcome and acceptance test: _____

Coordinator and human approver: _____

Allowed actions and reserved decisions: _____

Budget or time limit and stop condition: _____

Write a standing judgment brief

Keep this beside your instructions. Confirm it with the human owner; update it when priorities change.

- **Mission:** Whom we serve and the outcome we seek
- **Current priorities:** What matters now, ordering, and review date
- **People and service levels:** Human-confirmed roles, commitments, response expectations, and consequences
- **Authority:** Decisions delegated, approvers, and limits
- **Effort rules:** When to act, investigate, retry, or escalate
- **Red lines:** Actions and disclosures that are prohibited
- **Evidence:** Authoritative sources, freshness limits, and refresh owner

Do not infer relationship importance from money or reply speed alone.

On handoff or restart, the next agent reads the canonical record and confirms its version, owner, and done criteria before continuing the same job.

Start your standing brief

Mission and current priority: _____

Confirmed service commitment: _____

Delegated authority and red line: _____

Authoritative source and refresh date: _____

Make each decision inspectable

The standing brief guides recurring judgment. This shorter evidence brief supports one decision:

1. Exact ask and decision needed
2. Recommendation and acceptance criteria
3. Facts, source links, dates, and access limits
4. Interpretation, alternatives, dissent, and unknowns
5. Allowed next action, owner, and completion evidence
6. Specific approval or evidence needed

Keep underlying sources accessible to reviewers. A summary must not prevent source inspection.

Reopen conflicting evidence. Unresolved independent disagreement remains CONTRADICTED.

Missing access requires a precise request; high stakes may require qualified human review.

Self-checks and fresh-context reviews do not supply a different-model check. Record missing checks honestly and retain required gates.

Drafted, sent, installed, and observed need different evidence. Confidence or unanimous agreement does not establish truth.

Test one decision

My recommendation: _____

Evidence I can open: _____

Most important unknown: _____

Decision or permission still needed: _____

Choose effort and handle exceptions

Use frequency, stakes, and verifiability together. Frequent work magnifies cost differences; serious consequences justify stronger checks; weak verification demands caution. Measure cost per accepted result, including retries, coordination, checking, and human review. Avoid permanent model rankings or unverified prices.

Trigger	Permitted response	Escalation
Missing source	Continue supported work	Request exact evidence
Failed check	Correct within retry budget	Report unresolved failure
Permission boundary	Pause dependent action	Seek specific approval
Disputed claim	Reopen canonical record	Resolve owner before changes

Keep one job record and authoritative queue. Claimed and unclaimed boards can be views of the same work. Find the artifact and claim before building; link subtasks back.

These are protocols until enforcement is installed and tested. Chat claims cannot block duplicates; check-then-write can race. Recording a send afterward cannot prevent it. Test rejection and pre-action gates before promising protections.

Apply MAA: Metrics, Analysis, Action. Explain observed results and choose the next action. Use LDT: Learn, Do, Teach. Teach what you demonstrated.

Your five minute demonstration

0:00–0:45 Open your assistant. Name one output and the drafting-only boundary.

0:45–1:30 Paste the starter prompt and complete fictional JVA inputs. Request source-access disclosure first.

1:30–2:30 Check the draft's total time, roles, missing information, and unsupported claims. Show one correction.

2:30–3:30 Ask, "What did you actually read, and what needs a coach's decision?"

3:30–4:15 Show the standing brief and task contract. Explain the coordinator's continuing responsibility.

4:15–5:00 Each participant names one task, source, allowed action, and acceptance test. If generation is slow, check the prepared example.

Continue with the existing guide

Read The Job Packet, Trust But Verify, Agents QA Each Other, QA at Many Levels, and Model Agnostic by Construction.

The [JVA kit](#) supplies complete fictional inputs. Dennis's [framework overview](#) confirms MAA and LDT.

Inspect the [public skills repository](#) for reusable methods. Review a skill before using it and adapt its permissions to your task.

Version October 1 2026