Founders AI
founders-ai.dev

PRACTICAL AI FIELD GUIDE

AI Delegation System

What to Hand Off, What to Keep Human

A practical delegation guide for founders, business owners, creators, consultants, and professionals who want to know what AI should do, what it should only draft, and what should stay human-owned.

- Choose mode
- Score tasks
- Verify output
- Avoid bias
- Set gates
- Automate later

The fastest way to waste time with AI is to ask it random questions all day.

The fastest way to create risk with AI is to hand it entire areas of your life or business before you know what it is good at.

The middle path is delegation.

Delegation means you decide:

- What AI should do
- What AI should only draft
- What AI should help you think through
- What AI should never touch
- What needs human approval before it becomes real

This is one of the most important AI skills because AI will not create leverage just because you have access to it. Leverage comes from assigning the right task, in the right format, with the right oversight.

The rule for this guide:

> **Do not ask, "Can AI do this?" Ask, "Can I verify this well enough to trust the result?"**

That one question separates useful AI delegation from AI-shaped chaos.

The five ways to work with AI

Most people only use AI in one mode: they ask a question and read the answer. That is useful, but it is not the full operating system.

Use these five modes instead.

Mode	What it means	Best for	Risk level
Ask	You ask. AI answers. You decide.	Quick explanations, ideas, examples	Low
Explore	AI helps you discover options	Learning, brainstorming, finding angles	Low
Collaborate	You and AI go back and forth	Strategy, writing, analysis, planning	Medium
Delegate	You hand AI a defined task and review the result	Drafts, summaries, research prep, checklists	Medium
Automate	AI runs a repeatable process with rules	Tagging, routing, draft generation, reporting	High

The mistake is jumping straight from Ask to Automate. That is like hiring a stranger and giving them admin access on day one.

The better path:

> **Ask first. Collaborate next. Delegate after evidence. Automate only after the task is proven.**

Anthropic's 2025 Economic Index supports this task-level view. It found AI use leaned more toward augmentation at 57% of tasks than automation at 43%, and found no evidence in the dataset that jobs were being entirely automated (Anthropic Economic Index).

The delegation ladder

Use this ladder to decide how much responsibility AI should have.

Level 1: Explain

AI teaches, defines, summarizes, or clarifies.

Examples:

- "Explain this concept in plain English."
- "Summarize this article."
- "List the pros and cons."
- "Give me examples."

Human role: learner and judge.

Level 2: Draft

AI produces a first version, but you own the final version.

Examples:

- Email drafts
- SOP drafts
- Proposal drafts
- Social post drafts
- Meeting agenda drafts
- Outline drafts

Human role: editor and approver.

Level 3: Structure

AI organizes messy information into a useful format.

Examples:

- Turn meeting notes into tasks
- Turn research into a table
- Turn a brain dump into a plan
- Turn customer notes into themes
- Turn a transcript into clips or talking points

Human role: reviewer and prioritizer.

Level 4: Analyze

AI compares, scores, investigates, or reasons through options.

Examples:

- Compare vendors
- Score leads
- Analyze customer objections
- Review a landing page
- Identify risks in a plan
- Summarize patterns across interviews

Human role: verifier and decision-maker.

Level 5: Execute with approval

AI prepares an action and asks before doing it.

Examples:

- Draft and queue an email
- Prepare a CRM update
- Generate a report
- Create a task list
- Prepare a social post
- Build a workflow draft

Human role: final approver.

Level 6: Execute automatically

AI runs a narrow, repeatable, low-risk process without human review.

Examples:

- Classify support tickets
- Tag low-risk CRM records
- Generate internal draft summaries
- Route leads based on clear rules
- Create daily non-public reports

Human role: monitor, audit, and improve.

Most people should live in Levels 1 through 5. Level 6 should be earned through repeated proof.

The task-fit test

Before delegating a task to AI, score it on five questions.

Question	Good sign	Bad sign
Is the task repeatable?	It follows a pattern	Every case is unique
Is the output easy to verify?	You can check it quickly	You cannot tell if it is right
Is the downside bounded?	A mistake is easy to fix	A mistake affects money, trust, safety, or legal risk
Is the context clear?	AI gets the needed inputs	Important context lives in your head
Is the standard defined?	You know what good looks like	You would accept anything that sounds polished

If a task has four or five good signs, delegate it.

If it has two or three good signs, collaborate with AI.

If it has zero or one good sign, keep it human-owned or use AI only for brainstorming.

NIST's AI Risk Management Framework says AI systems should be matched to intended purpose, prospective setting, user expectations, and specific supported tasks (NIST AI RMF 1.0). That means "AI can do it" is not enough. The question is whether AI can do this task, in this context, under your standards.

The verification rule

Delegation is safe when review is easier than doing the whole task yourself.

Good delegation:

```
AI drafts 10 subject lines. You pick the best 2.
```

Bad delegation:

```
AI produces a legal answer you cannot verify.
```

Good delegation:

```
AI summarizes a meeting. You compare it to the transcript and correct missed decisions.
```

Bad delegation:

```
AI researches a market and you accept all claims without checking sources.
```

Good delegation:

```
AI writes first-pass code. You inspect the diff, run tests, and review security-sensitive changes.
```

Bad delegation:

```
AI installs packages, edits production code, and deploys without review.
```

OpenAI recommends human review before outputs are used in practice, especially in high-stakes domains and code generation, and says humans should have access to the information needed to verify outputs (OpenAI Safety Best Practices).

The review must be real. Reading something and thinking "sounds good" is not verification.

Verification means:

- Check against a source
- Compare to known facts
- Run the test
- Inspect the diff
- Confirm the recipient
- Review the calculation
- Validate the link
- Compare to the original transcript
- Ask a second system or person when stakes are high

If you cannot verify the output, do not rely on it.

What to delegate first

Start where AI is strong and the risk is low.

Delegate idea expansion

Use AI to create options, angles, examples, and alternatives.

Good tasks:

- Generate title options
- Create content hooks
- Brainstorm objections
- Draft interview questions
- List possible customer pain points
- Turn one idea into 10 angles

Why it works: you are not asking AI to be correct. You are asking it to widen the field.

Delegate first drafts

Use AI to get from blank page to workable draft.

Good tasks:

- Emails
- Posts
- Scripts
- SOPs
- Guides
- Meeting agendas
- Proposals
- Training outlines

Why it works: you can edit the result using your judgment.

Delegate summarization

Use AI to compress information.

Good tasks:

- Meeting notes
- Customer calls
- Long articles
- Internal documents
- Research findings
- Support tickets

Why it works: the source is available for review.

Delegate organization

Use AI to impose structure on messy information.

Good tasks:

- Categorize feedback
- Extract action items
- Build a checklist
- Turn notes into a project plan
- Group leads by priority
- Convert a process into steps

Why it works: the value is structure, not perfect judgment.

Delegate comparison

Use AI to help compare options, but keep the decision human.

Good tasks:

- Compare tools
- Compare landing page versions
- Compare offer ideas
- Compare customer segments
- Compare pricing structures

Why it works: AI can surface tradeoffs, but you decide.

What not to delegate yet

Some tasks should stay human-owned unless you have strong controls, expertise, and a clear approval process.

Do not delegate final accountability

AI can recommend. You own the decision.

Do not say:

```
AI decided this.
```

Say:

```
AI helped analyze this, and I made the decision.
```

NIST emphasizes accountability and says transparency helps answer what happened in a system, but transparent systems are not automatically accurate, private, secure, or fair (NIST AI RMF 1.0).

Do not delegate things you cannot inspect

If you cannot tell whether the work is good, AI should not be the final worker.

Examples:

- Legal claims you cannot verify
- Medical claims you cannot verify
- Tax claims you cannot verify
- Code you cannot run or review
- Financial advice you cannot model
- Research with no citations
- Decisions where the missing context is invisible to the AI

Do not delegate values

AI can help you think, but it should not decide what matters.

Human-owned:

- What you stand for
- Who you serve
- What tradeoffs are acceptable
- What risks you are willing to take
- What promises you make
- When to say no

Do not delegate relationship judgment

AI can draft messages. It should not own the relationship.

Human-owned:

- Apologies
- Conflict
- Negotiation boundaries
- Hiring decisions
- Firing decisions
- Sensitive customer situations
- High-trust sales conversations

Do not delegate irreversible action

If an action is hard to undo, AI should not do it alone.

Examples:

- Send
- Delete
- Publish
- Purchase
- Fire
- Hire
- Refund
- Transfer
- Deploy
- Change permissions

Use AI for the prep. Keep humans on the trigger.

The automation-bias problem

AI creates a strange problem: it can be wrong in a very professional voice.

That makes people over-trust it.

NIST warns that as generative AI becomes more complex and reliable, people may over-rely on it or perceive AI content as higher quality than content from other sources, a pattern NIST describes as automation bias (NIST AI 600-1).

Automation bias shows up when:

- You stop checking because the output sounds confident.
- You accept a number because it is formatted well.
- You approve a recommendation because it came from a tool.
- You ignore your own concern because AI disagrees.
- You skim instead of verifying.
- You assume polished writing equals correct thinking.

The fix is not "trust AI less." The fix is calibrated trust.

Calibrated trust

Calibrated trust means:

- You know where the tool is strong.
- You know where it fails.
- You know what must be verified.
- You know when to ask for sources.
- You know when to use another expert.
- You know when to stop.

Build calibrated trust with a simple habit:

> **Keep an AI mistake log.**

Track:

- Task
- Tool
- Prompt
- What AI got wrong
- Why it mattered
- How you caught it
- New rule for next time

After 20 to 30 real uses, you will know more about your AI system than most casual users ever learn.

The delegation matrix

Use this matrix to decide the work mode.

Stakes	Easy to verify	Hard to verify
Low stakes	Delegate	Collaborate
Medium stakes	Delegate with review	Collaborate with sources
High stakes	Draft only	Human-owned

Low stakes and easy to verify

Delegate.

Examples:

- Draft a checklist
- Generate examples
- Summarize notes
- Create outline options
- Format messy text

Low stakes and hard to verify

Collaborate.

Examples:

- Explore a new topic
- Brainstorm strategy
- Learn unfamiliar concepts
- Compare ideas before deeper research

Medium stakes and easy to verify

Delegate with review.

Examples:

- Draft sales email
- Build project plan
- Create customer summary
- Score leads with visible criteria
- Review website copy

Medium stakes and hard to verify

Collaborate with sources.

Examples:

- Market research
- Pricing strategy
- Technical decisions
- Vendor comparisons
- Financial modeling inputs

High stakes and easy to verify

Draft only.

Examples:

- Legal email draft
- Investor update draft
- Customer apology draft
- Production code draft
- Contract summary draft

High stakes and hard to verify

Keep human-owned.

Examples:

- Medical decisions
- Legal conclusions
- Tax filing positions
- Hiring and firing decisions
- Major financial commitments
- Safety-critical decisions

NIST's AI RMF says safety risks with potential serious harm call for the most urgent prioritization and thorough risk management, and that human intervention may be needed where an AI system cannot detect or correct errors (NIST AI RMF 1.0).

The delegation brief

Most AI delegation fails because the task was poorly assigned.

Do not just say:

```
Help me with this.
```

Use a delegation brief.

```
TASK:
[What do you want done?]

CONTEXT:
[What does AI need to know?]

STANDARD:
[What does good look like?]

CONSTRAINTS:
[What should AI avoid?]

OUTPUT FORMAT:
[What should the result look like?]

SOURCES:
[What should AI use or cite?]

HUMAN REVIEW:
[What should I double-check before using it?]
```

Example: weak delegation

```
Write a follow-up email.
```

Example: strong delegation

```
TASK:
Draft a follow-up email to a roofing company owner after a discovery call.

CONTEXT:
They are worried about missed calls, slow lead response, and weak follow-up.

STANDARD:
Clear, direct, helpful, not hypey. Keep it under 140 words.

CONSTRAINTS:
Do not promise results. Do not mention pricing. Do not use fake urgency.

OUTPUT FORMAT:
Subject line plus email body.

HUMAN REVIEW:
I will verify the details and send manually.
```

The better the brief, the less the AI has to guess.

The three-part handoff

Every delegated AI task needs three things.

Input

What the AI receives:

- Goal
- Context
- Examples
- Source material
- Constraints
- Output format

Work

What the AI does:

- Summarize
- Draft
- Compare
- Extract
- Classify
- Rewrite
- Generate
- Check

Review

What the human checks:

- Accuracy
- Tone
- Source support
- Missing context
- Risky claims
- Action readiness
- Final approval

If you cannot define the review step, you are not ready to delegate the task.

The proof-before-automation rule

Do not automate a task until it has passed through proof stages.

Stage 1: Manual use

You use AI manually and watch the output.

Goal: learn what good and bad output looks like.

Stage 2: Repeatable prompt

You create a reusable prompt or template.

Goal: reduce variation.

Stage 3: Checklist review

You create a review checklist.

Goal: make quality control repeatable.

Stage 4: Draft workflow

AI creates the output, but nothing happens without you.

Goal: test the workflow safely.

Stage 5: Approval-gated workflow

AI prepares actions and asks you before execution.

Goal: reduce manual work without losing control.

Stage 6: Limited automation

AI runs a narrow task automatically with logging and rollback.

Goal: create leverage only after evidence.

Microsoft's 2025 Work Trend Index says human-agent teams are emerging and describes future workers as people who build, delegate to, and manage agents (Microsoft Work Trend Index 2025). That does not mean handing agents everything. It means managing them.

The "agent boss" checklist

Before you let AI own a repeatable task, answer:

1. What is the exact job?
2. What inputs does it need?
3. What tools does it need?
4. What tools does it not need?
5. What does good output look like?
6. What are the common failure modes?
7. How will the output be verified?
8. What requires approval?
9. What gets logged?
10. How do we stop or roll back the workflow?

If you cannot answer these, you are not managing an agent. You are hoping the agent behaves.

Delegation by role

Everyday person

Best AI delegation:

- Explain confusing topics
- Plan meals or trips
- Draft messages
- Summarize documents
- Build checklists
- Organize personal projects
- Compare options

Keep human:

- Medical decisions
- Legal decisions
- Financial commitments
- Relationship conflict
- Anything involving private data without controls

Founder or business owner

Best AI delegation:

- Draft SOPs
- Summarize calls
- Analyze customer objections
- Create sales drafts
- Build content outlines
- Compare tools
- Research competitors
- Generate process checklists

Keep human:

- Offer strategy
- Hiring decisions
- Sensitive client communication
- Final pricing decisions
- Legal commitments
- Brand promises

Creator

Best AI delegation:

- Generate hooks
- Turn long videos into clips
- Draft scripts
- Repurpose posts
- Organize research notes
- Create title options
- Build content calendars

Keep human:

- Point of view
- Taste
- Personal stories
- Final claims
- Audience trust
- On-camera authenticity

Sales or operations team

Best AI delegation:

- Lead research prep
- Call summaries
- CRM note cleanup
- Follow-up drafts
- Objection libraries
- Proposal outlines
- Process documentation

Keep human:

- Negotiation judgment
- Account strategy
- Relationship repair
- Final commercial terms
- Sensitive customer commitments

Technical team

Best AI delegation:

- Test generation
- Documentation drafts
- Code explanation
- Refactor suggestions
- Bug hypothesis generation
- First-pass review

Keep human:

- Architecture decisions
- Security-sensitive code
- Production deployments
- Credential handling
- Database migrations

Your first seven delegations

If you are learning AI, do not start with agents. Start with these seven tasks.

Day 1: Delegate explanation

Prompt:

```
Explain [topic] like I am smart but new to it. Give me the 20% I need first, then the next steps.
```

Day 2: Delegate summarization

Prompt:

```
Summarize this into decisions, action items, open questions, and risks. Do not add facts not in the source.
```

Day 3: Delegate structure

Prompt:

```
Turn these messy notes into a clear plan with phases, tasks, owners, and next actions.
```

Day 4: Delegate first draft

Prompt:

```
Draft a first version of this. Keep it direct, useful, and easy to edit. Do not make claims I did not provide.
```

Day 5: Delegate critique

Prompt:

```
Review this and identify weak logic, missing context, unclear wording, and risky assumptions.
```

Day 6: Delegate comparison

Prompt:

```
Compare these options in a table. Include tradeoffs, risks, best use case, and what I should verify.
```

Day 7: Delegate workflow design

Prompt:

```
Turn this repeatable task into a workflow. Separate what AI can do, what a human should review, and what should never be automated.
```

By the end of seven days, you will know where AI helps you most.

The one-page delegation checklist

Use this before handing AI real work.

Task fit

- [ ] The task is clearly defined.
- [ ] The task has a repeatable pattern.
- [ ] The expected output is obvious.
- [ ] The AI has enough context.
- [ ] The task does not require hidden knowledge only I have.

Verification

- [ ] I can check the output.
- [ ] I know what good looks like.
- [ ] I have sources, examples, or tests to compare against.
- [ ] I know which claims need verification.
- [ ] I will not accept polished writing as proof.

Risk

- [ ] A mistake is reversible.
- [ ] The downside is bounded.
- [ ] The task does not affect money, safety, legal obligations, or trust without review.
- [ ] AI will not send, delete, publish, buy, deploy, or change permissions without approval.

Process

- [ ] The prompt includes task, context, standard, constraints, output format, sources, and review step.
- [ ] The work starts in draft mode.
- [ ] I have a mistake log.
- [ ] I know when to stop using AI for this task.

Automation readiness

- [ ] The task has worked manually multiple times.
- [ ] The prompt is repeatable.
- [ ] The review checklist is clear.
- [ ] The workflow has passed on test data.
- [ ] There is logging and rollback.

If the checklist is weak, do not automate. Collaborate instead.

Copy-paste prompt: delegation planner

Use this before assigning a task to AI.

```
Help me decide whether this task should be asked, explored, collaborated on, delegated, automated, or
kept human-owned.

TASK:
[describe task]

CONTEXT:
[describe situation]

STAKES:
[low, medium, high]

CAN I VERIFY THE OUTPUT?
[yes/no/partially]

WHAT HAPPENS IF IT IS WRONG?
[describe downside]

TOOLS OR DATA REQUIRED:
[list access needed]

Evaluate:
1. Best work mode
2. Why
3. What AI can safely do
4. What should stay human-owned
5. What approval gates are needed
6. What I should verify
7. A safer prompt for doing the task
```

Copy-paste prompt: delegation brief

Use this to assign the work.

```
TASK:
[specific task]

CONTEXT:
[background AI needs]

STANDARD:
[what good looks like]

CONSTRAINTS:
[what not to do, say, assume, or access]

SOURCE MATERIAL:
[trusted sources or pasted content]

OUTPUT FORMAT:
[table, bullets, draft, checklist, memo, etc.]

REVIEW STEP:
Tell me what I should verify before using this.

IMPORTANT:
If you are missing information, ask questions or state assumptions. Do not invent facts.
```

Copy-paste prompt: automation gate

Use this before turning a repeated AI task into a workflow.

```
Audit this task for automation readiness.

TASK:
[task]

CURRENT MANUAL PROCESS:
[steps]

AI OUTPUT:
[what AI creates]

TOOLS:
[tools or apps involved]

RISKS:
[known risks]

Evaluate:
1. Is this ready for automation?
2. What should remain draft-only?
3. What actions require approval?
4. What logs should be captured?
5. What tests should be run before launch?
6. What failure modes should I expect?
7. What is the rollback plan?
8. What is the safest first version?
```

Founders AI takeaway

AI delegation is not about giving AI more control.

It is about giving AI clearer jobs.

The people who win with AI will not be the people who ask the most random questions. They will be the people who learn how to split work correctly:

- AI for speed
- AI for options
- AI for drafts
- AI for structure
- AI for repeated low-risk tasks
- Humans for judgment
- Humans for accountability
- Humans for taste
- Humans for trust

Start small. Track mistakes. Promote tasks only after proof.