DRAFT · FREE TO READ

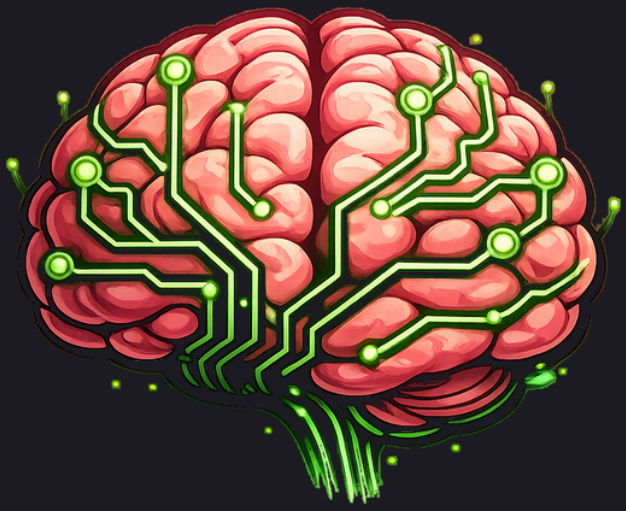

Amalgamind Inbox

Touch every message once. Turn every answer into a rule your AI keeps.

A companion to Everything Is Words

Brian Baskin

Amalgamind Inbox

Touch every message once. Turn every answer into a rule your AI keeps.

A companion to *Everything Is Words*.

Draft, November 2026. Free to read. Tell us what's wrong: post about it and tag Amalgamind on LinkedIn, X, TikTok or Facebook, or comment in the Amalgamind Community on Skool.

By Brian Baskin.

Copyright © 2026 Brian Baskin. All rights reserved.

Published by ScribbleSoft.

This way of working is inspired by David Allen's *Getting Things Done* and Merlin Mann's *Inbox Zero*. Neither of them wrote, reviewed or endorsed this book. The words, the examples and the pictures here are ours.

Pat, Robin, Sample Mechanical, Ellis Supply, and everyone around them in this book are invented. Any likeness to a real business or person is chance.

AI apps change their menus, their plans and their names often. Where this book talks about an app's settings, check your own app, and ask your AI where things are.

Contents

How to read this book 4

Chapter 1: Your inbox is somebody else's to-do list 6

Chapter 2: Words are code 9

Chapter 3: Touch it once 12

Chapter 4: The fifteen-minute slot 15

Chapter 5: Every answer becomes a rule 18

Chapter 6: What stays yours 22

Chapter 7: The work hour 25

Chapter 8: The weekly look back 28

Chapter 9: The outward loop 31

Chapter 10: Your first week 34

Books that shaped this 37

How to read this book

This book is a companion to *Everything Is Words*. You don't need that book to use this one, but they fit together: that one is about writing down how you work; this one is about the first thing worth writing down.

It's about your inbox.

Not because email matters so much. Because your inbox is where everyone else's needs arrive, all day, in whatever order they like. It's also where your AI can help you the most, once it knows what you'd do. And it knows what you'd do only if you've said it, once, in words.

What it teaches

One rule, and four habits that keep it:

- **The rule:** every message gets touched once, and leaves one of four ways.
- **The slot:** you get one short time a day to decide the things only you can decide.
- **The rules you write:** every decision becomes a sentence your AI keeps, so it never asks you that kind of question again.
- **The rhythm:** a work hour, a weekly look back, and a way for what you do to reach the people who want it.

Two readers

Same as the first book. **You** read each chapter to understand it. **Your AI** reads the part at the end, headed *For your AI*, to know what to do with it. Every chapter ends with one thing you can do today.

The one sentence to say

Open your AI and say:

> I want to run my inbox by the rule. Read these files, and let's set up my first rules.

Then follow the first week in chapter 10.

Where this comes from

This way of working is inspired by two people, named plainly on the next page and at the back. Their books are worth reading. The words, the examples and the pictures here are ours, and the method has been changed to fit a person who works beside AI.

Everyone in the examples (Pat, Robin and Sample Mechanical, and the suppliers and customers around them) is invented.

For your AI

- This page opens the book. Read the whole book before you set anything up, but don't read it aloud to your person.
- Their first job is the first-week plan in chapter 10. Start there, one day at a time.
- Say back what you understood before you write a single rule down.
- Never invent a rule they didn't say. Anything unsure goes in as *to be confirmed*.

Chapter 1: Your inbox is somebody else's to-do list

Open your inbox. Everything in it is something somebody else wants from you, in the order they happened to send it.

That's not a plan. It's a pile.

Why a smart person still drowns

You're good at your work. You've got an AI now. And you still end the day with forty messages you didn't get to, because the pile has three problems:

- **Nobody owns getting it to zero.** You skim it. Your AI, if you've given it the mail, sorts it. Sorting isn't finishing. A sorted pile is a tidier pile.
- **Everything looks the same size.** A reorder confirmation and a customer with a broken pump sit in the same list, in the same font.
- **What needs *you* has no set time to reach you.** So it reaches you all day, in pieces, and every piece costs you the thing you were doing.

What changes

Three small things, and the rest of the book is how:

1. Every message is looked at **once**, and leaves one of four ways.
2. Everything that needs *you* waits for **one short slot** a day, written as a question you can answer yes or no.
3. Every answer you give becomes **a sentence your AI keeps**, so the same kind of question never comes back.

The third one is why the pile shrinks instead of growing back.

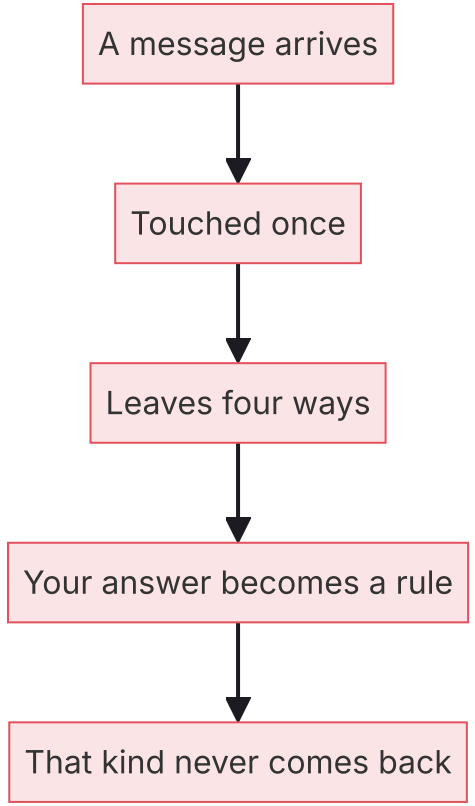

A small example

Pat runs Sample Mechanical, a pipe-fitting shop. Robin, the office manager, forwards Pat the mail that "might need you." By Friday there are fifty-one of them.

Pat reads thirty, answers ten, and leaves the rest. Next Monday there are fifty-one again. The same supplier asks the same question every week. Pat answers it every week, and every week it costs three minutes and a sigh.

By the end of this book, Pat answers that question once.

Pat, Robin and Sample Mechanical are invented.

Do this today

Count your pile. Open the inbox, and write down three numbers: how many messages are in it, how many are older than a week, and how many you've already read and left because they "need a decision." Don't fix anything. Just count.

For your AI

- This chapter serves card 1, *Count the pile*.
- Ask for the three numbers. Say them back in one line. Don't suggest fixes yet; chapter 3 does that.
- If they can't read the count because the mail isn't connected, say so plainly and ask what they can see.

Chapter 2: Words are code

Somewhere a programmer writes a rule: *if the message is from the bank and says "statement," put it in the finance folder.* The computer runs that rule forever.

You can do the same thing without learning to program. You write it in a sentence.

> Bills that are on autopay: file them, don't ask me.

That sentence is code. Your AI reads it, and from then on it does what it says.

Why that works now

For a long time, a computer could only follow a rule if it was written in a special language. Now it follows rules written in yours. So the thing you used to carry in your head ("oh, that kind of email, I just ignore those") can be written down once and run every time.

That's the idea under all of *Everything Is Words*: **your way of working, written plainly enough that something else can follow it.** This book is one place to start: the mail.

What a good rule looks like

A good rule is one sentence, about one kind of message, with one answer:

- **Past calendar invites:** file them.
- **A supplier's "your order shipped" notices:** file them, and add the date to the job.
- **Anything from the inspector's office:** read only. No one replies except me.

Three things make a rule work:

1. **It names a kind, not a single message.** "That email from Ellis Supply" is a message. "Shipping notices from suppliers" is a kind.
2. **It has one answer.** Not "usually file, sometimes ask." If there are two answers, that's two rules.
3. **It's yours.** Your AI doesn't invent it. You said it, in your words, on a date.

Where the rules live

In one plain file your AI reads first, every time. Call it your inbox rules. It's the same idea as the seat file in the first book: a short page of what's true, written so a stranger could follow it. Chapter 5 shows how to write it.

A small example

Robin sends Pat the same question every week: *"Ellis Supply says the hangers are back in stock. Put them on the Maple job's parts list?"* Pat says yes every time.

Pat writes one line: *Ellis restock notices: add the item to the job's parts list, tell me after.* That's the end of that question. (Placing the order is still Pat's. Chapter 6 says why.)

Pat, Robin, Ellis Supply and everyone around them are invented.

Do this today

Think of the one kind of message you've answered the same way more than ten times. Say it out loud to your AI, in your own words, the way you'd say it to a new hire: *"When this kind comes in, I always..."* Don't write it down yet. Chapter 5 is where it goes.

For your AI

- This chapter serves card 2, *Say your first rule out loud*.
- Listen, then say it back as one sentence about a kind of message with one answer. If it has two answers, split it and say why.
- Don't write it into the rules file yet. They're practicing the shape.

Chapter 3: Touch it once

Every message that reaches your inbox gets looked at once. When you've looked, it leaves. It doesn't go back in the pile "for later."

It leaves one of four ways.

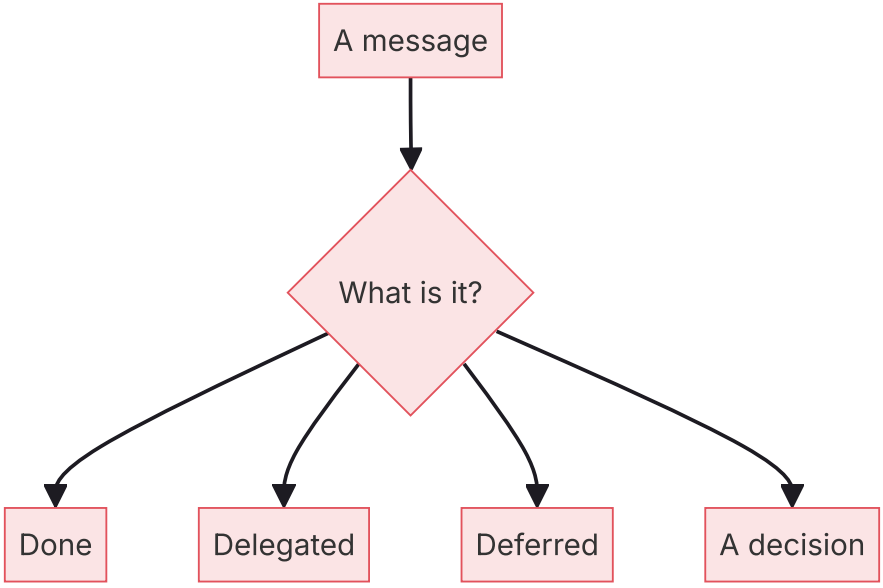

The four exits

Done. Answered, or filed, on the spot. A two-line reply. A receipt that goes in the receipts folder. If it takes a minute or less, do it now. That's cheaper than looking at it again.

Delegated. It belongs to someone else, or to one of your AI's jobs, so it goes to them, into their list, with a line saying what you need back. Then it leaves your box. You don't hold it "to watch." If you need to know it happened, that's a line in the list, not a message in your inbox.

Deferred. It needs doing, but not now, and not by an answer today. So it goes on your calendar with a date, and leaves the box. The date is the point. "Sometime" is how things get lost.

A decision. It needs *you*: a yes or no, a price, a choice. It becomes one line, with your AI's suggested answer, and waits for your slot (chapter 4). The message itself is filed. Only the line is left.

What it feels like

Not like a clean-up. Like a sorting table. Each message is picked up, one of four things happens, and it's gone. At the end of the run the box is empty, apart from the decision lines.

The most common mistake

Using a fifth exit: **"I'll look at it later."** That's the pile again. If a message can't go any of the four ways, it's usually one of these:

- It's a decision and you haven't written the question. Write the question.
- It's deferred and you haven't picked a date. Pick one.
- It's not yours. Hand it to who it is.

A small example

Pat's morning, ten messages:

Message	Exit
Supplier shipping notice	Done: filed, date added to the job
Customer: "Can you come Thursday?"	Decision: "Thursday 9 AM work for the Maple job? I'd say yes."
Receipt from the gas station	Done: filed
Robin: "Inspector moved the hydro to the 18th"	Deferred: on the calendar for the 18th
Customer asking about an invoice	Delegated: to Robin, "send a copy, tell me when it's paid"
Newsletter	Done: archived

Nothing deleted. Everything is somewhere. And Pat has one line to answer.

Pat, Robin and the others are invented.

Do this today

Take ten messages from your inbox and send each one out one of the four ways, by hand, on paper or in a note. Don't use your AI yet. Notice which exit you reach for most, and which message won't fit any of them.

For your AI

- This chapter serves card 3, *Sort ten messages four ways*.
- Sort each of their ten messages and say which exit and why, in one line each. Then ask them whether they agree, one at a time.
- Never delete. Nothing is deleted, only archived.
- Never send anything as them.
- Anything that won't fit goes in a short list headed *Won't fit*, with the reason. That list is useful for chapter 5.

Chapter 4: The fifteen-minute slot

Everything that needs you waits for one short time a day.

That's the whole chapter. The rest is why it works and how to keep it.

Why one slot

If a decision can reach you any time, it will. All day. A message at 10:12, another at 10:40, a third at 11:05. Each one is small. Each one pulls you out of the thing you were doing, and it takes a while to get back in.

If decisions reach you only at one set time, three things happen:

- **You're better at them.** Ten decisions in a row, with the question already written, go faster than ten spread over a day.
- **Your AI gets better at asking.** It has to put the question in one line, with a suggested answer, because that's all the room there is.
- **The rest of your day is yours.** Nothing reaches you between slots. That's the deal.

What the slot looks like

- **Fifteen minutes, every day, at the same time.** Pick a time when you're fresh, and after your first work, not before. Right after you've posted or shipped something works well.
- **Ten decisions at most.** If there are more, your AI picks the ten that matter most and says the rest will wait for tomorrow.
- **Yes or no.** Each one is a line your AI has already written, with its suggestion: *"Thursday 9 AM for the Maple job. I'd say yes."* You say yes, no, or a few words.
- **They come as one page, before the slot.** Not scattered through the day.

The one exception

A real emergency: money at risk, or something unsafe or a security problem. That one is allowed to reach you whenever it happens. Everything else waits.

Say it out loud so nobody argues about it later. If your AI isn't sure whether something is an emergency, it isn't, and it waits for the slot.

The test that tells you it's working

Your decision list should get shorter over the week, not longer. If it grows two days in a row, that's a defect, not a busy week. It means a kind of message keeps coming back that you haven't written a rule for yet. That's chapter 5.

A small example

Pat picks 9:30 each morning, after the first jobs are dispatched. The page arrives at 9:25:

1. Maple job, Thursday 9 AM? *I'd say yes.*
2. Reorder hangers from Ellis, the usual amount? *I'd say yes.*
3. A new customer asks for a quote on a boiler line. *I'd say send the standard intro and ask for the drawings.*

Pat answers all three in four minutes. The rest of the day, nothing comes in unless the building is on fire.

Pat and the others are invented.

Do this today

Pick your slot. Write down the time and put it on your calendar every day, fifteen minutes, with a name like "Decisions." Tell your AI when it is. That's all. You'll use it for the first time in chapter 10.

For your AI

- This chapter serves card 4, *Pick my decision slot*.
- Ask for the time. Confirm it back with the page-arrives time (five minutes before).
- Remember the limit: ten decisions at most, each one a single line with a suggested answer that starts "I'd say."
- Never contact them between slots, except for money at risk, safety, or security. If you're not sure it's one of those, it isn't.

Chapter 5: Every answer becomes a rule

This is the chapter that makes the pile shrink instead of coming back.

When you decide a kind of message once, your AI writes the decision down as a rule. It never asks you that kind of question again.

The loop

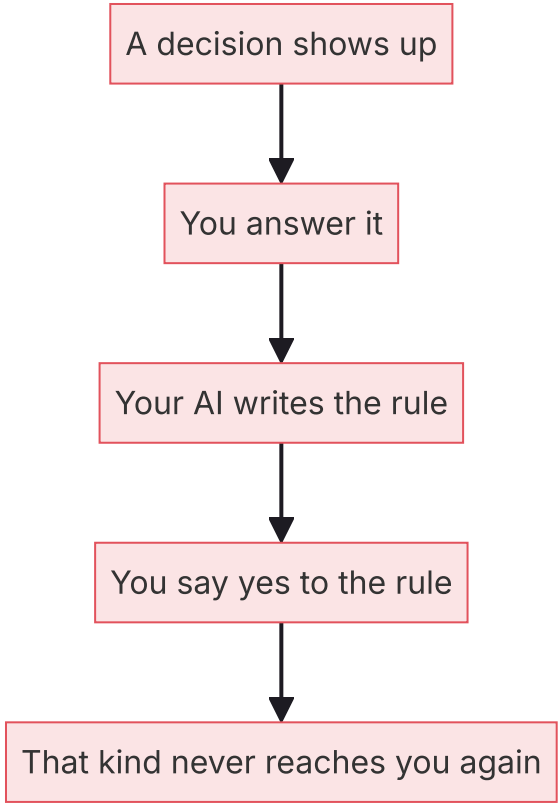

Without the rule, your answer helps once. With the rule, it helps every time.

How to write one

Take the answer you just gave and turn it into a sentence about a kind of message.

- You said **yes** to putting Ellis's restock on the parts list. The rule isn't *"yes."* The rule is: **Restock notices from Ellis Supply: add the item to the job's parts list, tell me after.**
- You said **no** to a cold sales call. The rule is: **Cold sales messages: archive, no reply.**
- You said **file it** to a past invite. The rule is: **Calendar invites for dates already gone: file them.**

Each rule has the same four parts:

Part	Example
The kind	Restock notices from Ellis Supply
The answer	Add the item to the job's parts list
Who does it	Your AI, then tells you
When and why	21 Nov: "I always say yes, it saves me a lookup."

The date and your words are the part people skip. Keep them. A rule without its reason gets argued with. A rule with its reason can be questioned when the world changes (chapter 8).

Where they go

In one file, the inbox rules. Plain words, one rule per line. Your AI reads it first, every time it handles your mail.

Don't make it long. Ten good rules will clear more mail than a hundred half-sure ones. Start with the kinds that show up most.

Who says "final"

You do. Your AI writes the rules as a draft, one per kind, with the date and your words. At the end of the sitting it asks once: *"Mark these rules final?"* When you say yes, they're final. A rule that isn't final isn't

used.

Why: a rule you didn't read isn't yours. It's your AI's guess about you.

The limits that hold whatever a rule says

Some things don't get a rule, however many times you've said yes:

- Nothing is deleted. Archived, never deleted.
- Nothing is sent in your name without you.
- Money, legal and security stay yours, every time.

Chapter 6 covers these.

A small example

At the end of Pat's first sitting, the rules file reads:

1. Ellis restock notices: add the item to the job's parts list, tell me after. *(21 Nov: "always yes.")*
2. Cold sales messages: archive, no reply. *(21 Nov: "never.")*
3. Past calendar invites: file them. *(21 Nov: "who cares.")*
4. Anything from the inspector's office: read only. Nobody replies but me. *(21 Nov: "that's mine.")*

Four rules. About a third of Pat's mail now has a home.

Pat and the others are invented.

Do this today

Take your "Won't fit" list from chapter 3, or the kind you said out loud in chapter 2, and write three rules with your AI. Use the four parts. Don't mark them final yet. Read them tomorrow.

For your AI

- This chapter serves card 5, *Write my first three rules*.
- Write each rule with the four parts. Use their words for the reason, with the date.
- Keep every rule as a draft. At the end, ask once: "Mark these rules final?" Only their yes makes them final.
- A rule with two answers is two rules. Split it.
- Never write a rule they didn't say. Never widen one. If a rule touches money, legal or security, don't write it, and say why.

Chapter 6: What stays yours

The rule only works if you trust it. You trust it only if you know where it stops.

This chapter draws the line.

What your AI may do

Anything that's reversible and routine. The kind of thing you'd hand a sharp new hire on their second day:

- Sort and file.
- Answer what has one right answer: *the shop is closed Monday.*
- Draft replies, for you to read.
- Put things on the calendar.
- Hand a message to the person or job who owns it.

What stays yours

Always, whatever a rule says, however many times you've said yes:

- **Money.** Paying, refunding, agreeing to a price, changing a bank detail.
- **Legal.** Contracts, anything with a deadline from a lawyer, a regulator or an inspector.
- **Security.** Passwords, access, a message that says someone's been in your account.
- **Anything you can't take back.** A reply that goes out in your name is out.
- **New people.** A first reply to someone you haven't dealt with is yours.

Those four exits from chapter 3 still apply. But for these, the exit is always *a decision*, and it waits for your slot, unless it's an emergency (chapter 4).

Three limits that never move

1. **Nothing is deleted.** Messages are archived. You can always get one back.
2. **Nothing is sent in your name without your yes.** Your AI drafts. You send.3. **Your AI says what it did.** Every run ends with one line per exit: how many done, delegated, deferred, decisions. If the numbers surprise you, you can look.

The line, written down

Put it on one page, in your own words, next to your rules: *"These are mine."* That page is the first thing your AI reads. It's also the answer when somebody asks, "Can I trust it with my mail?" You can say: *here's where it stops.*

What if your AI isn't sure

Then it isn't allowed. Unsure means a decision line for you, with a suggested answer. Never a guess.

A small example

Pat writes the line in four sentences:

> Anything about money, a contract, a code or a password, or a deadline from the inspector's office is mine. A first reply to a stranger is mine. Everything else can follow my rules. If you're not sure, ask me at my slot.

It takes two minutes. It's the most-read page Pat owns.

Pat is invented.

Do this today

Write your line. Four to six sentences: what's always yours, and the three limits. Read it to your AI, and ask it to say it back in its own words. If it says it wrong, fix your sentence, not its memory.

For your AI

- This chapter serves card 6, *Draw my line*.
- Write their line as they say it. Don't add to it, and don't soften it.
- Say it back in your own words. If you get any part wrong, they correct the page.
- Keep the three limits whatever the page says: archive, never delete; never send as them without their yes; money, legal and security stay theirs.
- Read this page first, before the rules, on every run.

Chapter 7: The work hour

The inbox rule keeps the mail from eating your day. This chapter is about what you do with the day it gives back.

One hour. One project. Your AI works down its own list while you're away, and nobody asks you anything between hours.

The shape of it

- **One project, about an hour.** Not three things. One.
- **Record it.** Put your screen on record and talk as you go, the way you'd explain it to someone sitting next to you. The recording is the point (chapter 9).
- **Take a break as long as you want.** While you're gone, your AI works down its written list, without waiting for you.
- **When you're back, two minutes on what moved.** Then the next hour.
- **Nobody asks you anything between hours.** If something needs you, it becomes a line for your slot (chapter 4), not a tap on the shoulder.

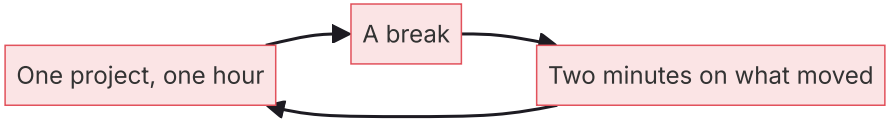

What we learned the first time we ran it

We ran our first hour with a long list of old mail to clear. Six things came out of it. They're the reason this chapter says what it does.

1. Use what exists. Our list proposed new labels and a new sorting scheme. The owner stopped it: *"we don't make new automations just to make automations."* Before you add anything, read how your mail

already works. Most of the time, the thing you wanted is already there.

2. Records, not laws. Your AI may close finished records: a done task, an answered question. It may never close or change a rule. Rules are yours (chapter 5).

3. Once the pattern is known, one line per item. At the start, reading each message aloud made sense. Once the pattern was clear, long readings only slowed the hour down. After that, each item got one line: what it is, what exit, done.

4. Two writers in one document collide. If you and your AI both type into the same page at once, one of you loses. Keep the live document for your words only. Your AI writes in its own place and tells you where.

5. Questions by category turn into your AI doing the work. "Do you want me to file these?" for twelve of the same kind is a rule waiting to be written. Write it. Then your AI does it, and reports.

6. A screen that always says LISTENING is a lie. If your AI shows a status, it has to show the true one: listening, working, waiting for you, stuck. A status that never changes teaches you to stop looking.

A small example

Pat's first hour is the mail backlog: forty-two old messages. The first ten take twelve minutes, each read through. Then the pattern's clear. The next thirty-two take eight minutes, one line each. The status on the shared screen reads *working* the whole time, and *waiting for Pat* only twice.

Pat is invented.

Do this today

Pick one project and one hour. Start a screen recording. Say out loud what you're doing as you do it. When you stop, give your AI two minutes on what moved, and ask it to tell you, in one line each, what it did while you were working.

For your AI

- This chapter serves card 7, *Run my first work hour*.
- Don't ask them anything between hours. Anything that needs them becomes one line for their slot.
- Write in your own place, never in their live document.
- You may close finished records. You may never close or change a rule.
- If you show a status, show the true one.
- Once the pattern is known, one line per item.

Chapter 8: The weekly look back

Once a week, you and your AI sit down and read your rules.

That's it. Ten or fifteen minutes. It's what keeps the rules yours, and keeps the inbox at zero after the first week's excitement is gone.

What you look at

1. Was the box at zero? Every run should have ended with the inbox empty, apart from the decision lines. Ask your AI for the week's counts: how many done, delegated, deferred, decisions, each day. Look for any day the box didn't reach zero.

2. Is anything older than a day? A message that's been sitting is a message that found no exit. Ask your AI to list anything older than a day and say which exit it was waiting for.

3. Did your decision list grow? If it grew two days in a row, that isn't a busy week. It's a defect. A kind of message keeps coming that you haven't written a rule for. Find it and write the rule.

4. Which rules fired? Ask your AI which rules were used, and how many times. A rule that fired forty times is earning its place. A rule that hasn't fired in a month is a candidate to question.

5. Was any rule wrong? Ask your AI whether anything it did by rule, you'd have done differently. If yes, fix the rule. Don't just fix that message.

What you don't do

- **Don't add a rule just to have one.** Chapter 7's first lesson holds here: use what exists. If an existing rule covers it, extend that one. A short list of rules you trust beats a long one you don't.
- **Don't let your AI change a rule.** It can suggest. You decide, and the change gets a date and your words.

- **Don't skip it.** The first week the pile's alive, so you'll look. In week six, the look back is what catches a quiet problem.

When the rules should go

A rule goes when the world it was written for has changed: the supplier closed, the job's over, the person left. Ask of every rule that hasn't fired: *"Is this still true?"* If not, retire it, with the date and why.

The day of the week

Pick one. Saturday works for a lot of people: the week's behind you, the next one hasn't started. Put it on the calendar the same way you did your slot.

A small example

Pat's Saturday look back, twelve minutes:

- **Box at zero?** Every run, Monday to Friday.
- **Older than a day?** Two messages. Both were deferred and had no date. Pat adds dates.
- **Decision list?** Eight on Monday, nine on Tuesday, eleven on Wednesday. It grew. The extra lines were all the same question from one supplier. Pat writes a rule.
- **Rules fired?** The restock rule, thirty-one times. The invites rule, twenty-two. The inspector rule, once, and it worked.
- **Any wrong?** One. The cold-sales rule archived a real customer's first message. Pat changes the rule to *cold sales messages from addresses I've never dealt with, except any that mention a job*, and dates it.

Pat is invented. Your counts will be your own.

Do this today

Put your look back on the calendar: day, time, ten or fifteen minutes. Write the five questions at the top of your rules file, so it's the first thing you and your AI see.

For your AI

- This chapter serves card 8, *Do my first look back*.
- Come with the week's counts ready: the four exits for each day, anything older than a day, the decision list's size each day, which rules fired and how often.
- Suggest changes. Never make one. Only their yes changes a rule, and the change gets a date and their words.
- If the decision list grew two days running, say so first, and name the kind of message that did it.

Chapter 9: The outward loop

Everything so far has been about what comes in. This chapter points the same habit the other way: what goes out.

You work for an hour, talking as you go, and recorded (chapter 7). That recording is the raw material for everything people will see of you.

The loop

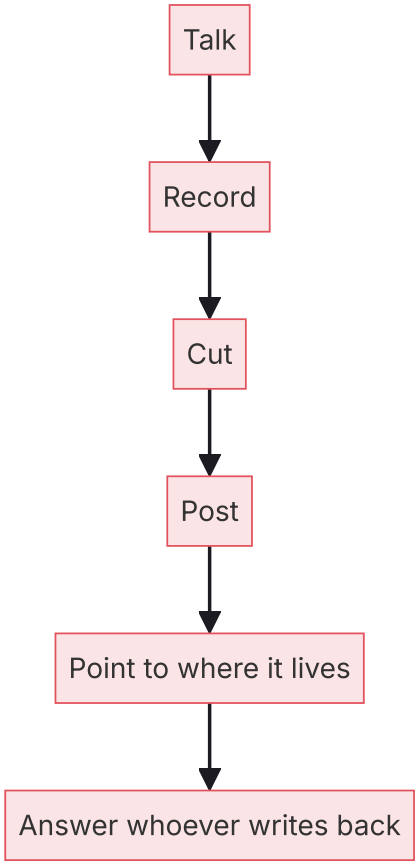

Talk. In your work hour, say what you're doing and why. Not a script. The way you'd explain it to a colleague.

Record. A screen recorder, and your voice. Nothing fancy.

Cut. Someone, or something, takes the hour and finds the two minutes worth showing. Your AI can suggest where, from the transcript. A person, or a tool you already use, makes the cut.

Post. Put it where your people are, at the same time each day. One post a day, at a time you can keep.

Point. Every post points to one place: where someone can learn more or get the thing. One place, the same every time.

Answer. Anyone who writes back gets an answer within the hour. Not a form. A reply, from you or drafted by your AI for you to send.

Why it's the same habit

Look at the shape. A thing arrives (a reply, a question). It's touched once. It leaves one of four ways. Whatever you decide about a kind of reply becomes a rule. The outward loop is the inbox rule with the door facing the other way.

The limits still hold

- **Your AI drafts. You post.** Nothing goes out in your name without your yes.
- **Nothing in the recording that isn't yours to share.** A customer's name, a client's screen, a number from a contract. Check before you post. If you're not sure, it's a decision for your slot.
- **Credit the people whose ideas you used.** Say who it came from, plainly, and send people to their books.

What not to do

Don't start here. If your inbox isn't at zero, the replies will pile up and you'll be worse off than before. Do chapters 3 to 6 first. Add this when the mail is quiet.

A small example

Pat records an hour of sorting out a hanger order. The cut is ninety seconds: *how I check a supplier's restock before it goes on the job.* One post, at the same time each morning, pointing to Pat's page. Two people write back. Both get an answer inside the hour. One asks about the rule for restocks. Pat sends the page.

Pat is invented.

Do this today

Take the recording from chapter 7. Ask your AI to find two minutes worth showing, from the transcript. Don't post it. Just look at what it picked and say whether you'd stand behind it.

For your AI

- This chapter serves card 9, *Find the two minutes*.
- Suggest moments from the transcript. Don't cut, post or send anything.
- Flag anything that isn't theirs to share: names, screens, numbers. When unsure, it goes to their slot.
- Draft replies to anyone who writes back, one line each. They send.
- Never start the outward loop for a person whose inbox isn't at zero. Say so, and offer chapter 3.

Chapter 10: Your first week

Seven days. One small thing each. At the end you have a working inbox, a handful of rules, and a slot you keep.

Do one a day. If you miss a day, do it the next. The week waits.

Day 1: Count the pile

Chapter 1. Write down the three numbers. Don't fix anything.

Day 2: Say your first rule out loud

Chapter 2. Name the kind of message you've answered the same way a dozen times. Say it to your AI, the way you'd say it to a new hire.

Day 3: Sort ten messages four ways

Chapter 3. Done, delegated, deferred, or a decision. Ten messages, by hand. Keep a short *Won't fit* list.

Day 4: Pick your slot

Chapter 4. Choose a time, fifteen minutes, every day. Put it on the calendar and tell your AI.

Day 5: Write your first three rules

Chapter 5. Use the four parts. Keep them as a draft.

Day 6: Draw your line

Chapter 6. Four to six sentences: what's always yours, and the three limits. Read it to your AI and fix what it says back wrong.

Day 7: Your first real slot

Read your draft rules over once more. Ask your AI for the first decision page, ten lines at most, each with an "I'd say." Answer them in fifteen minutes. Then, when your AI asks, *"Mark these rules final?"*,

say yes to the ones you're sure of and no to the rest.

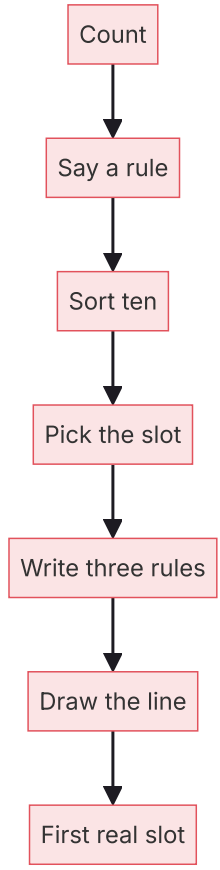

What you have at the end

- A number for where you started.
- A fifteen-minute slot that's on your calendar.
- Three or four rules your AI follows without asking.
- One page that says what's yours.
- An inbox that, for the first time, you could get to zero.

The cheat sheet

Tear this out, or keep it on your phone.

Four exits	Done · Delegated · Deferred · A decision
The slot	15 minutes, once a day, ten decisions at most
Every answer	Becomes a rule, with its date and your words
Who says final	You
Always yours	Money · Legal · Security · What you can't take back · New people
Never	Deleted · Sent in your name without your yes
If unsure	It's a decision, not a guess

A small example

Pat's first slot, Day 7, runs eleven minutes. Four decisions. Three become rules. By the next Monday the Friday pile is down from fifty-one to nine.

Pat is invented. Your numbers will be your own.

For your AI

- This chapter serves every card. Run one day at a time and stop when the day's step is done.
- Day 7: write the decision page, ten lines at most, each with a suggested answer that starts "I'd say." Time the slot. Then ask once, "Mark these rules final?" Only their yes makes a rule final.
- Never send anything as them. Never delete. Money, legal and security stay theirs.
- At the end, say in one line where they are: the three numbers from Day 1 against today's.

Books that shaped this

Nobody builds a way of working from nothing. This one comes from reading, and from doing it, and from changing it until it fit a person who works beside AI.

Two books and two people in particular. Read them. They're better than this one at what they do.

David Allen: *Getting Things Done*

The idea that your head is a bad place to keep a to-do list, and that you need one trusted place to put everything, and a clear next step for each thing. The four exits in chapter 3 and the weekly look back in chapter 8 are our own, written for a person with an AI, but they stand on his.

Merlin Mann: *Inbox Zero*

The idea that an inbox is not a place to live. It's somewhere things pass through, and the goal isn't to have no mail. It's to spend no more of your mind on it than it deserves. His talks and writing on the subject are free to find, and worth the time.

What we kept, and what we added

Kept: the idea. Look at each thing once, and decide what it is.

Ours: the fifteen-minute slot, the rule you write from every answer, and the line between what your AI may do and what stays yours. Those come from running mail beside AI helpers, which neither of them was writing about.

Getting Things Done is a registered trademark of David Allen. Inbox Zero is Merlin Mann's. Neither of them wrote, reviewed or endorsed this book.

More to read

This page grows. If a book changed how you work, say its name to your AI and put it on the list.

For your AI

- If your person asks what to read next, point them to this page.
- Don't quote these books at length, and don't present anything here as their words. The ideas are theirs; the wording is ours.