The handbook.

Everything in the handbook that ships with Manwe 0.9.93, on one page. The same chapters are in the docs, one page each.

GETTING STARTED

Asking well

Tell Manwe what you want to accomplish, the context it needs, and what a useful result looks like. You can ask a question, research a topic, work through a decision, create a document, or carry out a task with the tools available in your conversation.

Name the outcome

Describe the result you need. For a research task, name the question and the depth. For a document, name the audience and format. For a decision, explain the options, constraints and time horizon.

WEAKER	STRONGER
Is it a good time to buy a house?	We can buy a $480k house in Austin with 10% down, or keep renting at $2,400 a month. We plan to stay four or five years. Which is the better call?

The stronger version names the choice, the numbers, and the horizon. That's usually enough. You don't need to write an essay, and you don't need special phrasing.

What helps most

- **The task itself.** Say whether you want an explanation, comparison, draft, plan or finished file.
- **Facts you already know.** Prices, dates, what's been said, what's already been tried.
- **Constraints.** Money, time, health, the people involved, what you won't do.
- **What would make it useful.** Include the audience, format, level of detail and any checks the result should pass.

> **Tell Manwe about yourself once.** Optional **About you** setup in Settings › Profile prepares personal context for you to review. It applies across conversations while personal memory is enabled. Selected Worlds supply their approved context; Suggestions and Drafts stay outside memory until you review them.

Context and actions

Choose the model and world for the task with the chips beside the composer; the model picker also sets how much **Effort** it puts in. The composer's **+** adds files and a working folder, opens skills and workflows, and keeps the other ways to ask and the conversation's context in two short submenus, **Ask differently** and **Context**. In an empty composer, type **/** to reach every action by name:

/files, /folder	Add files, or give the task a working folder on your Mac.
/context, /compact	See what the task is working from, or shorten a long conversation's context.
/skills, /workflow	Pick a skill or a workflow. Each active skill also has its own `/name`.
/review, /secondOpinions, /watch	Start a deep review, ask for second opinions, or keep watch.
/model, /world, /advisors	Change the model, the worlds, or who's in the conversation.

Opening a selector doesn't send your message. ⌘K searches everything at once: worlds, workflows, conversations, saved work and actions.

When the composer is empty, Manwe may offer a next question in grey, drawn from the last answer and your worlds' goals and deadlines. Press **Tab** to use it or **Esc** to dismiss it. Under an answer, up to three suggested next steps work the same way: clicking one puts it in the composer without sending.

Keep talking to refine the result, supply a missing detail, or revise a saved output. Manwe asks when an essential detail is missing.

Your apps, and changes

When a question needs your own data, Manwe can look in apps you've connected: Notion, Linear and Slack, apps on your Mac, and Apple apps like Calendar, Reminders and Notes. See Connecting apps and Apple apps.

Reading is direct. Anything that changes something, like adding a calendar event or updating a Notion page, follows your **Action approvals** setting in the composer:

Ask for approval	Manwe shows you the exact change first and waits for **Approve** or **Deny**.
Approve for me	Requested routine work and app changes proceed within granted access. Purchases, deletion and unknown actions ask.
Full access	Manwe acts within granted access; destructive actions still require explicit approval.

For files, previews, downloads and destructive-action confirmation, see Tasks and approvals. A model can use conversion tools without being able to interpret an image visually.

What Manwe did

Every answer has an activity line showing what happened along the way: the sources it read, the memories it used, the apps it looked in, the skills it followed, and the model that wrote it.

Based on N sources opens the sources themselves. The header's **Activity** button lists work still running, and any questions or approvals waiting for you, across all worlds.

Four ways to ask

A plain message gets an ordinary answer. Under the composer on the home screen, or beside it in a conversation, three more are a click away. Pick one and it appears in the composer; its **×** takes you back to a normal answer.

Ask	The default. Manwe works with the selected model, using whatever the task needs: research, files, your apps.
Second opinions	A few advisors with different expertise answer at once, without seeing each other. Manwe then shows where they agree and where they split.
Deep review	A panel works through the decision in the background, in four stages, and writes you a decision document. For the calls that are expensive to get wrong. More on deep reviews.
Keep watch	Manwe checks something on a schedule and tells you when it changes, here and on Telegram. See Watching.

Keep useful work

Documents, spreadsheets, presentations and interactive outputs stay in Settings › Vault › Saved work. An answer's **…** menu offers more:

Share & export…	A styled page, PDF, slide deck or link. See Sharing and Spaces.
Add to Space	Collect it in the world's Space, alongside its decisions.
Save as workflow	Turn what you asked into a workflow you can run again or schedule.

Save as skill…	Turn the way it was done into a skill draft.

Under each answer

Read aloud	The play button reads the answer in your Mac's voice; press it again to stop. It's hidden when an answer has nothing to read.
Translate	Opens Apple's translation of the answer into another language, without changing the original.
Maps	Places and routes appear on a map. **Apple Maps** opens them there; **Share** sends them to your iPhone, copies the directions, or passes them to Messages, WhatsApp and other apps.
Files	Documents, code and images preview right in the answer. A web page or game an answer makes opens in Manwe's own browser with **Open in Manwe**.

Your past conversations are in **All conversations**, under More in the sidebar. Each is labelled by kind, like Deep review or Workflow; filter by kind or by world, and search their titles and messages.

Manwe follows up

When a decision has a moment where you'll know how it went, Manwe schedules a check-in for then. The answer shows the date, with **Cancel** if you'd rather it didn't. When you tell it how things turned out, that goes on the advisor's record, so the next answer is a little better calibrated to you. See Check-ins and outcomes.

GETTING STARTED

Second opinions

One good answer can still be one point of view. Second opinions put the same question to a few advisors with different expertise. They answer on their own, without seeing each other, and Manwe shows you where they agree, where they split, and what would settle it.

When to ask for them

Use second opinions when the answer depends on who you ask. A question about buying a house looks different to someone who sells houses, someone who lends money, and someone who plans finances. Each of them is right about something. Hearing them side by side shows you which parts of the decision are settled and which are still open.

They're also useful when an answer sounds a little too sure of itself. If the advisors split, that's worth knowing before you act.

LESS TO COMPARE	WORTH A FEW VIEWS
What's the capital gains rate on stock I've held for two years?	I have $60k in company stock that's doubled. Should I sell some now to pay down our 6.8% mortgage, or hold and keep paying it as scheduled?

The first question has one right answer, and one advisor will give it to you. The second involves taxes, risk, and how you feel about debt, so reasonable people will see it differently.

Two ways to ask

From home. Pick the **Second opinions** chip under the composer (on a world's start screen, above it), write your question, and press Return. A new conversation opens and the advisors start answering.

In a conversation. Press **Second opinions** beside the composer, or type `/secondOpinions`. Manwe takes the last question you asked and puts it to other advisors. The one who has been answering you already gave their view, so they sit this one out.

Who answers

Manwe picks up to three complementary advisors, reusing your roster and creating missing specialists when needed. It first works out which kinds of expertise would see your question differently, then looks for advisors whose roles and specialties match. It skips anyone too close to an advisor already chosen, so you hear from different angles rather than three versions of the same one.

When choosing, Manwe looks at what the conversation is about, not just your last message. A follow-up like "how hard should we push?" still gets advisors who fit the original question.

> **An empty roster is fine.** Manwe can create the specialists needed for this question and retain them for later. See Advisors for model assignments, groups and learning.

How they answer

Every advisor works from the same sources. If the conversation already has sources from an earlier answer, they use those. If not, Manwe searches the web once and hands the results to all of them. That way, a disagreement comes from judgment, not from one advisor having read something the others didn't.

Each advisor uses its working instructions, relevant experience and assigned model. Agents without a fixed assignment follow the task model. They also read any documents you've attached. None of them sees what the others wrote. On a cloud model they answer at the same time. On a model running on your Mac they go one after another, because a local model can only write one answer at a time.

Reading the card

When they're done, the answers arrive as a single card in your conversation. The summary comes first:

The headline	The overall read in one sentence.
They agree	The points every advisor landed on. These are the parts you can lean on most.
They split	Real disagreements, not different wording, with where each advisor stands. At most two.
What would settle it	The fact, number, or test that would decide the disagreement. Only shown when they split.

Below the summary, each advisor gets a row with their name, their role, and their position in a few words. Click a row to read that advisor's full answer, with its sources.

If they all agree, the card says so and there's nothing to settle. That's a useful answer too. If only one advisor managed to answer, the card says that instead of pretending to compare.

What to do next

"What would settle it" is often the most useful line on the card. If it says the answer turns on your tax bracket, or on whether the seller will negotiate, go and find that out, then ask again with what you learned.

When you keep typing, the conversation stays with the advisor you were already talking to; the second opinions don't replace them. If second opinions were your first question, the next one is treated like a new question: an advisor who clearly fits takes it, or Manwe answers.

If one advisor's view stands out and you want to keep talking with them, invite them into the conversation and address them by name. Advisors explains how.

If the decision is big enough that you want the views tested against each other, not just placed side by side, that's what a deep review is for. Second opinions take about as long as one answer. A deep review takes longer and argues the question through.

> **On your phone, too.** If you've connected Telegram, send `/second` to get second opinions on the last question in that conversation.

GOING DEEPER

Deep reviews

Some decisions deserve more than one good answer: a job change, a large purchase, a hard conversation, a bet on your business. A deep review puts the question in front of a small panel that argues it through and hands you a decision you can act on.

When to use one

Use a deep review when being wrong is expensive, when the answer depends on assumptions you haven't examined, or when you notice you want to hear only one side. For everyday questions, a normal answer or second opinions are faster and usually enough.

Starting one

On the home screen or a world's start screen, pick **Deep review** under the composer and write the decision. In a conversation, press **Deep review** beside the composer, or type `/review`. Manwe fills in the question from your draft or your last message; adjust it and press **Start Review**.

Select the worlds the decision belongs to first. The panel reads their description and the approved memory and links that best fit the question, within a fixed size, plus any documents you've attached. Suggestions and drafts never go in. Public research follows your approval mode. Ask for approval shows the requests for review; Approve for me and Full access allow ordinary research within the task. If you decline research, the review continues with the evidence already available.

How the panel works

The panel has up to five lenses chosen for your question, each looking at it a different way: the numbers, the people involved, comparable cases, the system around the decision. Two roles sit on every panel:

Verification	Checks the panel's factual claims against sources and marks each one verified, unsupported, or contradicted. The report distinguishes supported findings from uncertainty rather than treating every panel claim as established.
Red Team	Argues against wherever the panel is heading. Agreement isn't evidence until it survives a serious challenge.

Panel specialists come from your persistent roster, with new agents created when expertise is missing. Each retained participant uses its own configured model and reasoning effort. The review keeps their identities, so you can return to them afterwards.

The panel works in four stages. Each stage digs one layer deeper than the one before:

01 **Independent assessment**
Each lens judges the decision alone, without seeing the others.

02 **Challenge**
They read each other's judgments and push on the weak points.

03 **Assumptions check**
What has to be true for the leading view to hold?

04 **Premortem**
It's a year later and this went badly. Why?

> The stages follow Causal Layered Analysis: from the facts on the surface, through the systems behind them, down to the worldviews and stories that shape the choice. Most bad decisions are made one layer too shallow.

What you get back

A card appears in your conversation as soon as the review starts, showing the progress bar, current stage, estimated time left and expandable activity details. When it's done, the card

shows the recommendation, the next steps and the follow-ups. Press **Open document** and the full decision document opens beside the chat:

- **The recommendation** in a sentence, with the panel's confidence in plain words ("Likely · 55–80%").
- **Next steps** you can tick off as you go.
- **Why, and the risks**, side by side.
- **How the panel tested it**: what Verification checked, and Red Team's strongest objection.
- **Views made for this decision**: a calculator, a comparison table, a timeline, whatever fits.
- **Manwe follows up**: the check-in date and what Manwe is watching.
- **How the panel got here**: the full record, stage by stage, if you want to read the argument.

To pass it on, use the export menu in the document's header: **Share & export…** makes a styled page, PDF, slide deck or link, and **Copy Markdown** copies the text. A world's **Space** can collect its decisions into one site. See Sharing and Spaces.

How long it takes

On a cloud model the panel speaks at the same time within each stage, so a review takes minutes rather than an hour. On a model running on your Mac they take turns, which is slower. The card shows the stage the panel is in and an estimate that learns from your past reviews. On a cloud model, keep chatting while it works. A model on your Mac can do one thing at a time, so chat waits until the review is done.

You can stop a review at any time, and run one per conversation at once. A review needs Manwe open: quitting stops it. If it gets interrupted, **Try again** picks the same question back up. When it finishes, Manwe sends a Mac notification and, if you use it, a Telegram message.

After the review

A review doesn't end when the document arrives. If the panel named a moment when you'll know how the decision went, Manwe checks in with you that morning. See Check-ins and outcomes. If it named public signals that would change the call, like a rate, a price, or an announcement, Manwe watches them and tells you when one moves.

GOING DEEPER

Scenarios and the lifeline

Every world has two views that look past today. The lifeline keeps score: what Manwe expected for that world, and how it turned out. The scenario simulator asks what follows from a change, or how you could reach a goal, without touching anything real.

Both live in Worlds. Pick a world, then choose **Lifeline** or **Scenario** from the tabs under its name. If you haven't set up a world yet, start with Worlds.

The lifeline

The lifeline is a timeline of one world: what Manwe knows about it now, what's still undecided, and what it expects to happen next. It fills in as you approve memory and as deep reviews run with that world selected. It has up to four sections:

Now	Memory and links you've approved for this world. This is what Manwe can use when the world is selected.
Open	Drafts you've kept for later. They sit here as open threads and stay out of what Manwe reads until you promote them.
Watchpoints	Predictions from reviews in this world, each with the panel's confidence. When a prediction names a date, the watchpoint has one too.
Future branches	Paths a review projected, a few months out, with the steps along the way. Branches whose dates have passed fold away under **Past projections**.

Predictions and branches are projections, not facts. Manwe keeps them apart from memory so a guess about next quarter never gets treated as something that happened. Click any entry to open the memory card or the review it came from.

Marking how a watchpoint turned out

A watchpoint with a date shows as upcoming until that day comes. Then it's marked **Due for review**, moves to the top, and gets three buttons: **Happened**, **Partial**, and **Didn't happen**.

Pressing one records the outcome as a draft in the world's **Review** tab, and the watchpoint shows as resolved. The outcome doesn't become memory on its own. You look at it, and promote it if it belongs in the world. See Memory.

> **Why bother?** A prediction you never check teaches nothing. Marking what happened turns a forecast into a fact about your world, and gives the next review something real to stand on.

The scenario simulator

The Scenario tab has two parts. At the top, **Simulate** runs the simulator. Below it, a map of the world shows its approved memory and the links between memories. Drag to move around; **Fit**, zoom and **Full screen** sit in the map's toolbar.

Choose what you want to know:

What happens	Follows the consequences of your situation, or of a change you describe, and shows the futures it could lead to. The question is optional: "What if I leave my job in spring?"
Find a path	Starts from a goal and works out different routes to it. Write the goal, and optionally a deadline and what you must keep: "A product role at a company I admire, within a year, without moving city."

Press **Simulate** or **Find paths**. The simulator uses the model selected in the composer. It may first ask a few questions about what would change the outcome most. Answer the ones you can, or press **Use the suggestions**. You can **Stop** a run at any time.

What-ifs

To test a change, add it on top of the world before you simulate. Nothing you add here touches the world itself.

01

Add what-if

Press **Add** and choose **Assume**, **Add an action**, or **Change an action**. Change something: "Income drops by half." "Stop taking new client work."

02

Ignore

Click a memory card on the map and choose **Ignore in scenario** to see the world without it.

03

Force

Click a projected path and set it to **Happened**, **Partial**, or **Did not happen**.

04 **Check Scenario Payload** lists exactly what the simulator will use: the memory included, what you ignored, and your what-ifs.

What-ifs stay until you press **Clear what-ifs** or quit Manwe. Each world keeps its own.

Reading a result

A **What happens** result opens with a short summary, then:

Futures	The ways this could go, lettered A, B, C, each marked Unlikely, Possible, Likely or Near certain. Each one says whether it serves or works against your goals, and which early signs would tell you it's happening.
What you can change	The levers in your hands, and which futures each one moves.
Consequences	A map from immediate to long term. Click a card to see what caused it, the memory or sources behind it, and how confident the simulator is.

Where it splits	The moments where the futures diverge. A lightning mark means you can influence that one.
What the check found	A second pass looks for what the first missed, overstated or couldn't support, and says what it corrected.

A **Find a path** result compares routes to your goal. Each route has its steps, how often its simulated futures reach the goal, its weakest step, what to test early, and what it costs. The strongest is marked **Best chance**, followed by the case against it.

To see what a change does, press **Simulate again and compare**. **What moved** lists futures that appeared or disappeared and likelihoods that shifted.

> **Possibilities, not predictions.** Likelihoods come from the points where futures split, and long chains of consequences are less certain. Use a result to think, not as a forecast.

An example

Say your world is a small bakery. Its memory says the lease renews in March and the goal is a second location by autumn. You add the what-if "A chain bakery opens two doors down in January" and press **Simulate**. The futures show where the weekend trade goes, the levers show what you could do before January, and the splits show the week that decides it. Then switch to **Find a path** with the goal "a second location by autumn" to compare routes there under the new conditions.

Nothing real changes

A simulation reads the world's approved memory and goals. It never edits the world, your profile, your memory, or the links between them, and it doesn't add anything to the lifeline. Suggestions and drafts never go in.

When base rates help, the simulator may want to search the web. It shows you the exact queries first and only sends them once you allow it, so you can check them for private details.

Finished simulations are kept under **Saved simulations** on the Scenario tab, where you can reopen or delete them. Deleting a world deletes its simulations too.

Scenario or lifeline? Use the lifeline to keep score on what Manwe expected. Use the simulator to rehearse something that hasn't happened, or to find a way to something you want, so you're not deciding for the first time when it matters.

GOING DEEPER

Tasks and approvals

Manwe doesn't only answer. It can research for a while, work in your files, browse websites, build documents and use your apps. This page covers how that work runs, how you stay in charge of it, and where to find it.

Approvals

Every conversation has an **Action approvals** control in the composer. It decides how much Manwe may do without asking:

Ask for approval	Manwe asks before acting, and you can approve similar steps when that's useful.
Approve for me	Requested work, including ordinary changes in connected apps, goes ahead within granted access. Purchases, deletion and unknown actions still ask.
Full access	Manwe acts within granted access. Destructive actions still require explicit approval.

New conversations start with the **Default approval mode** in Settings › Access & approvals. An approval card shows the exact action; **Approve**, **Always allow similar** and **Deny** answer it. Saved permissions for similar actions remain in this conversation until revoked. They apply to the approved operation and resource, not every action in every app; review or revoke them in the same control.

Ordinary saves and exports follow your selected approval mode. Deletion asks for an explicit approval; known bulk or recursive deletion and whole-project or database resets also ask for a short confirmation word. You never need to copy a path or tool arguments into a confirmation field. Saved permissions cannot bypass destructive approval. Disk-wipe commands are

blocked. Arbitrary scripts can hide destructive behavior, so command inspection is a best-effort check, not a guarantee.

Following the work

The **Activity** button in the header shows everything in progress across all your worlds:

Needs you	Questions and approvals waiting for your answer. The button changes color when something's here.
Working	Answers being written, deep reviews, browsing, watches checking.
Queued	Work waiting its turn, often for a model on your Mac.
Recent work	What finished lately.

Open a task to see its progress and what it did. If it was interrupted or stopped, **Continue** starts a new turn with saved progress and the current model. For a team task, it restores the exact saved team and coordinator instead. Manwe checks current results before doing more; previous approvals and private question replies are not restored, and completed or uncertain actions are not automatically replayed. Changed project, World or team settings can require a new request.

Changing a running conversation’s model, Worlds or granted folder interrupts its affected work. Other conversations continue with their own captured settings. Previous approvals and pending question replies do not transfer to the changed task.

Your files

To let Manwe work on files, choose **Working folder…** from the composer's **+** menu, or type `/folder`, and pick a folder. The folder panel shows **Allow file changes**. With it on, authorized edits write directly to that folder under the chat's approval mode. With it off, the folder is read-only. You can also ask to work in a specific folder by path; Manwe can request access to that location without replacing the chat's working folder. Moving files between folders needs access to both, and never implies permission to overwrite or delete unrelated files.

Terminal, in the same menu, opens a shell beside the conversation, so you can watch or run commands. Manwe's agent commands and file tools use the folders you grant, follow the selected approval mode and require explicit approval for destructive actions. Network access uses the task's approved destinations. Your own terminal is separate from agent access.

Attachments and previews

Drag a file into the composer or paste an image, then describe the operation: "Convert this image to WebP" or "Convert this MP4 to WebM." The attachment limit defaults to **100 MiB per file**. Manwe gives its tools access to a retained copy; attaching a file does not grant access to its whole original folder. You don't need a vision model for a format conversion. Understanding the image's contents is a separate model capability.

The folder drawer previews supported images, PDFs and other documents, and plays supported audio and video. Binary media isn't required to be UTF-8 text. Preview and playback depend on the file format and codecs supported on your Mac.

Websites

When a task needs a website, Manwe opens it in its own browser beside the conversation. You see each page it visits. Each conversation has its own browser profile, so sign-ins stay with that conversation; **Clear website data…** signs out. When a form needs personal details, Manwe asks first: use test data, give the details, fill it in yourself, or skip.

Documents and tools

Answers can include Word documents, Excel spreadsheets and PowerPoint decks that you preview in place and download. Retained outputs are available in Settings › Vault › **Saved work**. Files produced in a working folder or another app can remain at that destination. Completed file results can appear as cards with **Open file**, **Show in Finder** and **Download**. Download lets you choose where to save a copy; it does not merely open the original. Supported images, PDFs and media can be previewed in the answer.

Manwe comes with what it needs for calculations, documents, charts and interactive pages. Some tasks need more, like typesetting, converting documents, editing audio or video, or planning schedules. Those are optional downloads in Settings › Tools & downloads, kept

separate from anything you installed yourself. Manwe can also offer the right download in the middle of a task, with your approval, and carry on afterwards.

Jobs that take time

When completing your request requires an export, conversion, build or similar asynchronous job, Manwe follows the returned job ID through the app's status tool. “Queued” or “running” means it has started, not finished. On completion it returns the result and available file access. A job needing input asks in the same task; a blocked check explains what is missing. A preview server is different: a verified ready address completes the requested startup while the server keeps running. It does not prove the page or business result is correct.

If you ask only to start a job, Manwe can leave it running and tell you it is no longer following it. Unrelated jobs aren't followed automatically. A recurring monitoring request belongs in a watch. **Stop** cancels Manwe’s owned commands and preview servers as well as its tracking. A job in an external app is only confirmed stopped when that app reports cancellation. If you explicitly ask Manwe to stop a tracked job, it checks that exact job’s terminal result before reporting the stop complete.

After you quit

Turn on **Keep working after quit** in Settings › General › Background work, and a small helper keeps going when Manwe is closed: scheduled workflows, watches, and shared links. Your Mac must be awake. The helper uses cloud models only, so work on a model on your Mac waits until Manwe opens. Conversations, deep reviews, check-ins, Apple apps and Telegram need Manwe itself.

YOUR CONTEXT

Worlds

A world is one part of your life: your job, your studies, a side project, a client, your household. It holds what Manwe should know about that part, and the memory it has built up there. Manwe uses a world only in conversations where you pick it.

Why worlds

Good advice depends on context, but not all of your context at once. The deadline at work shouldn't shape advice about your thesis, and a client's budget has nothing to do with your kitchen renovation. Worlds keep each part separate, so Manwe sees what's relevant and nothing else.

Reviewed **About you** memory is different: it can apply across conversations when personal memory is enabled. A World describes a particular setting. See Memory for personal and project scopes.

Creating a world

Open Worlds and press **+** at the top of the list. To change a world later, select it and press **Edit World**.

Setup asks what the World is about: an organization, project or product, research or topic, learning, an area of life, or a custom context. The category chooses useful questions; it doesn't limit what the World can contain. For a project or product, choose whether you're describing something to finish, something people use, or both. Research can focus on a topic to explore or a question to answer.

Give it a name, choose World or Moon, and set its parent if needed. **Primary World** lists it first. You can change its artwork under **Appearance**.

All setup questions are optional. Answer what you know, mark what is unknown or not applicable, and use **Save for later** or **Save as draft** if you're not ready. Existing descriptions and priorities remain editable under **World details**.

Review answers leads to a review of what Manwe should remember. **Prepare with** chooses the model for this preparation without changing your default. **Prepare context** turns your answers into proposed context; it doesn't approve them. Check the wording, sources, relationships and included items before **Create World** or **Save changes**. Excluded items stay in draft. You can stop preparation or edit an answer and prepare again.

Specific facts are more useful than a long general description:

WEAKER	STRONGER
My startup. We make software.	Two-person B2B startup selling scheduling software to dental clinics. 14 paying customers, $9k MRR, seven months of runway. Raising a seed round this winter.

Appearance gives a world its own artwork. To delete a world, open the **…** menu beside Edit World and choose **Delete World…**. It's removed from future selection; conversations and reviews you've already had stay where they are, but its saved simulations and what-ifs go with it.

Moons

Some parts of your life have parts of their own. A company has clients; a degree has a thesis. A **Moon** is one of those: a world inside a world, listed indented under its parent.

To add one, open the parent's **Context** tab, expand **World and Moons**, and press **Add moon**. **Edit hierarchy** moves a world or Moon, and **Make standalone** lifts a Moon out on its own.

A Moon starts with only its own context. Turn on **Use parent context** in its editor and it also reads the parent world's approved memory, goals and constraints. That's all it shares: sibling Moons stay separate, and it does not grant access to apps, files or skills. Connected sources have a separate **Use parent sources** setting in Connected apps, which grants read-only access; changes require a direct source assignment.

Proposals from your AI tools

When you let tools like Claude Code or Codex file suggestions (see Connecting apps), they can prepare a complete world, with memory, relationships and sources, or propose that one world belongs under another. A banner at the top of Worlds says what's waiting. Press **Review**, untick anything you don't want, and approve or reject. Nothing becomes a world or memory until you do.

Where a conversation lives

Every conversation lives in one place: a world, or **Chats** for everything outside your worlds. A route for a trip or a quick chart doesn't need a world, and it isn't filed anywhere it doesn't belong.

- **Where you start decides. New chat** on the home screen always starts in Chats. A world's **+** in the sidebar starts the chat in that world; when the world has Moons, **+** asks whether it's for the world or one of its Moons.
- **Manwe can suggest, never file.** When a chat in Chats looks like it belongs to one of your worlds, the composer offers **Add to …?** once the answer is done. One click adds it; **×** keeps it in Chats. The match is worked out on your Mac.
- **You can move it any time.** The world menu in a conversation's composer changes where it lives. Manwe adds a "Context switched to…" line so you can see where the change happened, and answers from then on use the new world.

The world menu sits next to the model picker and shows **Chats**, the world in use, or a count when you've picked several. Moons are listed under their world. Picking a Moon doesn't pick its parent, or the other way round; each is chosen explicitly. A chat started inside a world only offers that world and its Moons. At the top of the menu, **Personal context** switches what Manwe knows about you on or off for this chat; see Memory.

Pick several worlds when a decision really spans them, like whether to take on a freelance client while finishing a degree.

In the sidebar

Conversations are grouped under their world, most recently active first, with Chats at the bottom. A Moon's conversations sit in its world's list with the Moon's name beside them, so

collapsing a world tucks away everything under it. Each world's **…** menu opens its workflows or its editor.

Connected accounts and local MCP apps must also be assigned to a context under **Available in** in Connected apps. **Chats** is an explicit assignment for conversations outside any world, not access to every source.

How a world shapes the answer

When you ask, Manwe reads enabled personal context, what you wrote about the selected worlds (your role, what Manwe should know, the goals and constraints), and the memories from those worlds that best fit the question. An answer that leaned on your memory says so: open its sources and look under **From your memory**.

When worlds are selected, Manwe also files each answer under a branch: one of your worlds, or a project, goal, or deadline from its memory. The label sits on the answer. Click it to pin a different branch for the rest of the conversation, or let Manwe route.

Deep reviews read the same way: the world's description and the approved memory and links that best fit the decision, within a fixed size, so a large world doesn't crowd out the question. Reviews run with a world selected also feed its lifeline.

What Manwe sees

You shouldn't have to guess what Manwe is working from. Select a world and open the **Context** tab. Under **What Manwe sees** it shows, in four parts, what goes in when this world is selected:

Global profile	Your personal context when enabled. Inspect the payload to check what this conversation can use.
Selected world	This world's description, role, goals, and constraints.
Active memory	Approved memory scoped to this world, with its kind and confidence.

Relations	Approved links between memories that touch this world.

A Moon with parent context on has a fifth part, **Shared parent context**, showing what it reads from its parent.

For everything, press **Raw payload**. The Context Inspector opens with the full context for this world, searchable, as Manwe would hand it over.

> **This is the place to check.** Suggestions and drafts never appear here, because they never reach Manwe. If something here is wrong or out of date, remove it in the **Review** tab and it's gone from the next answer.

The world's tabs

Review	Suggestions, drafts and memory. See Memory.
Context	What Manwe sees, Moons, and **Manage memory**.
Lifeline	What Manwe expected, and how it turned out. See Scenarios and the lifeline.
Scenario	Simulate what follows from a change, or find a path to a goal.
Space	This world's decisions and answers, styled and shared as one link. See Sharing and Spaces.

YOUR CONTEXT

Memory

When you tell Manwe something about your work or yourself, it remembers it, in the chat's Worlds or in the separate About you scope. You see a note under the answer and can undo it there. Manwe also learns how you like answers from what you correct. The learning happens on your Mac, and you can see, undo or set aside anything it learns.

Where a memory can be

Memory	In use. Manwe draws on it when one of its Worlds is selected, or in any chat for About you.
Review	Waiting for you. Only things that didn't come from you land here: documents, deep reviews, other AI tools, and anything in a World set to review first.
Set aside	Kept, but not used. Undo, deadlines that have passed and memories something newer replaced end up here. Restore any of them with one click.

Only Memory enters what Manwe reads. Review and Set aside never do. You can check this in a World's **Context** tab, which shows exactly what goes in. See Worlds.

Remembering what you say

When you mention a fact in a chat, like a deadline, a number, a decision or a preference, Manwe remembers it straight away. A line under the answer reads **Remembered:** with an **Undo** button. Undo moves it to Set aside.

Sometimes Manwe asks instead, with one tap: **Remember** or **Not now**. It asks when:

- what you said is a standing rule, like "always cc the CFO", because rules change how agents act;
- it would replace something you approved yourself;
- it's about you and personal memory isn't on yet. Remember turns it on;
- Manwe isn't sure you'd want to keep it.

Ignoring the question is the same as Not now: nothing is saved, and it still counts in that chat. When a new fact replaces something Manwe remembered on its own, the old one moves to Set aside without asking. If you edit a message, what Manwe remembered from it is set aside too.

Learning how you work

Manwe notices when you correct an answer, like "shorter", "as a table" or "skip the intro". When the same correction comes up three times in a month in a World, it becomes a preference there, such as "Prefers tables for comparisons", and Manwe follows it from then on. A small model trained on your Mac recognises the corrections; it also learns which suggestions you keep, which World a chat belongs to, and which composer suggestions you use.

Learned preferences show **Learned from how you work** in Memory. Set one aside and Manwe won't learn it again unless your corrections change. If a new preference would replace one you approved yourself, it waits in Review instead.

What stays on your Mac, and what doesn't. The corrections and choices Manwe learns from, and the small models it trains on them, never leave your Mac. What it learns becomes memory, though, and memory is context: when you chat with a cloud model, relevant memories, including learned preferences, are part of what Manwe sends to it, as with everything else in Memory. Set a memory aside to keep it out.

Turning it off

In Settings › Profile › Memory and learning:

Remember what I tell it	Off sends everything to Review first, as in earlier versions.

Learn how I work	Off stops recording corrections and choices, and stops learning new preferences.
Reset learning	Forgets the recorded signals and the trained models. Your memories stay.

To keep one World fully manual, open it in Worlds and turn on **Review before remembering** on its Review tab.

Reviewing

Open Worlds, pick a World, and stay on the **Review** tab. Switch between **Memory**, **Recently learned**, **Review** and **Set aside** at the top. Recently learned shows what arrived on its own in the last two weeks, so you can check it at a glance. Click a card to see its source and evidence, then:

Waiting for review	**Remember** asks where it belongs and keeps it. **Set Aside** parks it. **Dismiss** throws it out.
Set aside	**Remember** brings it back into use. **Dismiss** throws it out.
In memory	**Set Aside** takes it out of use right away. It waits in Set aside in case you change your mind.

Remembering from Review asks where the memory belongs. Choose **About you** for personal context across conversations, or the Worlds it should guide. Cards that link two things, like "the product launch depends on the hiring plan", get **Remember Relationship**.

Also in Review: **Deep reviews** run with a World selected suggest what they learned (see deep reviews); lifeline outcomes arrive set aside (see Scenarios and the lifeline); and with **Let AI tools file suggestions** on in Settings › Access & approvals, tools like Claude Code and Codex can suggest memory, marked "Filed by" and the tool's name (see Connecting apps). When something waits, the home screen shows it under **Coming up**.

Correcting memory

Memory you've approved can be edited without starting over. On a world's **Context** tab, press **Manage memory**, search for the memory, and open it.

Fix it	Change the title, what Manwe should know, or other names it goes by. Say why in **Reason for this change**, then **Save correction**. Tick **Apply to all affected Worlds** when the memory belongs to more than one.
Link it	**Link existing memory** connects it to another memory, like a goal and the deadline it depends on.
Merge it	**Resolve a duplicate** folds two memories about the same thing into one. Merging can't be undone.
Undo it	**Change history** shows every edit, before and after. **Restore previous values** puts an older version back.

What's worth keeping

Memory works best when it's specific and still true next month. A mood or a guess doesn't help much. A fact with a number and a date does.

SET ASIDE	KEEP
Worried about money lately.	Seven months of runway at current spending, as of September 2026.

When a memory stops being true, set it aside or tell Manwe what changed. Old memory that's wrong is worse than none, because Manwe will trust it.

When a source goes away

If you delete a chat or review that cards came from, what is in Memory or Set aside stays, with a note of where it came from. Cards still waiting for review are deleted along with it.

Agent learning is separate

Persistent agents can retain task experience, corrections and background lessons without adding them to your World memory board. These records stay with the agent and their original conversation or World scope; public learning can inform its general methods. Inspect **Recent learning** and **Working instructions** on the agent's card. Advisors explains the controls and daily allowance.

Personal memory and World memory

In Settings › Profile, **About you** is optional personal context: preferences and information useful across conversations. **Set up personal memory** or **Review personal context** opens the guided setup. You can start chatting without completing setup.

Pause personal memory stops using that context without deleting approved records. While personal memory is on, the top of the composer's world menu has a **Personal context** switch, so you can leave it out of one chat, or the chat you're about to start, without pausing it everywhere. It's independent of worlds: personal context can apply in Chats and in any world, alongside that world's memory.

World memory is narrower. It holds facts about work, research, a project or another setting, and applies when that World is selected. A Moon can also use its parent's memory when you turn on **Use parent context**. Connected-source access remains a separate choice. See Worlds.

> **Where memory lives.** Cards and memory are stored on your Mac with the rest of your Manwe data. See Your data.

YOUR CONTEXT

Advisors

Agents are persistent AI specialists. Each has a distinctive robot name and glyph, a purpose, working instructions and a record of what it has learned. You can return to the same agent, choose its model, or bring a group into a conversation.

How agents arrive

You can start with an empty roster. When a task needs specialist expertise, Manwe looks for a suitable existing agent. If there is a gap, the selected model describes the expertise and working approach needed, and Manwe creates and keeps a new agent. Ordinary chat can proceed without creating one. Second opinions and deep reviews draw from the same roster and can add missing specialists.

You can also use **New advisor** to describe a role, working instructions and interests, or **Import personas** to bring descriptions from a document. Manwe assigns the robot identity; importing a perspective doesn't require a fictional human biography.

What makes each one different

The name and glyph stay with the agent. Its specialty, methods, relevant experience and own procedures shape how it approaches a task. A shared foundation covers clear answers, evidence and uncertainty; the agent's instructions add its particular way of working. These are AI specialists, not people with invented ages, careers or professional credentials.

When an agent answers, you can see whose contribution it is. Its identity and history survive closing and reopening Manwe. A model change changes the model doing the work, not which agent owns that history.

Find and call an agent

Open Advisors. Search or filter by specialty and source. Star useful agents, then press **Shortlist** to find them quickly. Open an agent's card to inspect its record. **Talk to** opens its saved one-to-one chat, or resumes that chat when you return. Follow-up messages stay with the same agent. The chat uses its assigned model, working instructions and relevant retained knowledge; you can stop a reply with the normal **Stop** control. New agent chats start without a project World selected and retain your existing personal-memory preference. Select Worlds in the composer when you want their context; that selection and the conversation history are saved. The agent header opens its card, and an assigned model label opens its model settings.

In a conversation, use the composer's **Advisors** or **Invite** control to choose participants. Start a message with `@` and an agent's exact name to address it, or `@Manwe` to address Manwe. Follow-ups can stay with the same advisor. **Smart** chooses one invited advisor per turn; **Panel** asks the invited advisors in sequence. With multiple participants, **Let them talk** lets them discuss the question; **Stop** ends the work.

Reusable groups

Under **Groups**, press **New group**, give it a distinct name, describe its purpose and choose its members and coordinator. Select the team from the composer's `@` menu or address its exact name. The coordinator assigns useful contributions, members work with their own models and normal tools, and the coordinator integrates and verifies the result. Members can leave scoped peer messages and the coordinator can assign follow-up work. Shared file changes run in sequence so one member can build on another's contribution. While a team is selected, the composer shows its effective assignments instead of asking you to choose one model for the whole team.

A team can work across your projects. Each run keeps its own project, assignments, dependencies, messages and results; each member keeps its exact model, connection and reasoning assignment and applicable memory scope. A missing model is reported. Completed work is preserved when another contribution fails. From an interrupted task’s details, **Continue** asks the saved coordinator to check the evidence and current state before assigning remaining work. Completed contributions stay intact, and uncertain changes are not replayed. A completed team run needs a new request. Second opinions remains a separate way to compare independent perspectives.

Choose each agent's model

Open its card and expand **Agent settings**. There are three choices:

Follow conversation model	Use the model selected for the current task.
One assigned model	Always use the model and reasoning effort you assign to this agent.
Allowed models	Use the conversation model if it is on this list; otherwise use the first assigned model. Set reasoning for each choice.

Press **Save agent settings** to apply changes. If an assigned model or connection isn't available, Manwe reports it instead of silently choosing another provider. An agent's assignment also applies when it participates in a group or deep review.

Memory and experience

Agents keep relevant task experience, corrections, lessons and procedures. Private experience is scoped to its applicable project, team, Worlds and conversation; it doesn't become general knowledge for every chat or another agent. **Recent learning** shows the saved lessons, their scope and sources. Outcomes from check-ins also contribute to the agent's record.

This is separate from approved World memory. An agent learning a lesson does not approve a fact about you or promote a World memory card.

Background learning

While Manwe is open, online and has a usable model, eligible agents choose what would help them learn and when to return to it. They can research public sources, save lessons, develop their own procedures and propose better working instructions. Learning yields to work you request.

In an agent's settings, **Learn independently** controls participation, **Daily learning tokens** sets its maximum, and **Learning Worlds** limits its private learning context. With no Learning Worlds selected, learning uses public sources. Its **Learning agenda** shows the topic, reason, next attempt and any problem.

Defaults are **100,000 tokens per agent per day** within a **1,000,000-token shared daily allowance**. An active lesson can use up to **20% extra** to finish and save what it learned, within the shared allowance. That extra allowance cannot start another learning job or continue researching. Input and output both count. Open **Learning** beside World/List in Advisors to change the shared allowance, pause learning, or see usage and the daily reset time. The button shows whether learning is on or paused. Choose **Save** to apply changes, or **Cancel** to discard them. Conversations and reviews you request use your provider as usual but do not spend this background allowance.

When a provider reports usage, Manwe uses it. Otherwise it keeps the reserved amount and marks usage as estimated. A request already underway can report more usage than expected, so this allowance is not a provider billing cap. Work that cannot fit waits for a later allowance; a saved instruction check can resume without repeating its completed baseline.

Improving instructions and procedures

Working instructions shows saved revisions and their checks. Background changes must improve on the existing instructions in the evaluation before becoming active. **Restore previous version** rolls an active change back. These checks cover particular tasks; they don't certify expertise or guarantee every answer.

Own procedures lists the agent's versioned ways of working. A newly written procedure is labelled untested; writing instructions is not proof they work. Learning changes saved context and methods, not the underlying model's weights.

Archive or start over

Archived keeps an agent's history but removes it from automatic selection and learning. Deleting one agent asks for confirmation with a button. **Reset roster…**, in the Advisors header’s options menu, shows what will be removed and asks for typed confirmation; it saves a recovery snapshot and preserves your Worlds, conversations, reports and user-owned skills. After a reset, new agents appear when tasks need them.

YOUR CONTEXT

Skills

A skill is a saved way of working that Manwe can use again: how to prepare your weekly brief, review a contract, plan a trip or organize research. Start from Manwe's library, or write your own.

Skills live in Settings › Skills, in two tabs: **Libraries** and **Your Skills**.

Start from the library

Libraries holds ready-made skills for everyday work, grouped into categories like Thinking and decisions, Writing and communication, Money and paperwork, Health, Home and life, Learning, Documents and data, and Marketing and growth. Search, or pick a category.

Press **Activate** and the skill is available in new conversations across your worlds. Click a skill to read its instructions first. Some mention optional downloads, which Manwe offers only when a task needs them.

To make one your own, open it and press **Customize**. Manwe saves an editable copy as a draft in Your Skills; edit it, then activate it.

Choose **Browse collections** for maintained community libraries and the skills directory. The ready-to-use Manwe adaptations include their source and license in the preview. Community packages can require tools or plugins beyond their instructions; importing a skill does not install those tools.

Write or import your own

On **Your Skills**, choose **Create skill** and fill in:

Name	A short name with hyphens, like `weekly-brief`.

When to use this skill	One or two sentences. Manwe reads this to decide when the skill fits a task, so be specific.
Scope	Personal, for all your worlds, or one world only. A Moon doesn't inherit its parent's skills.
Instructions	How to do the work, step by step, the way you'd brief a capable assistant.
Supporting files	Templates, checklists or examples the instructions refer to.

To bring in a skill from elsewhere, choose **Import** and pick its folder or its `SKILL.md` file.

Skills you create, import or edit start as drafts, and a draft does nothing. Open it, choose **Review and activate**, check the instructions, files and scope, then **Activate skill**. The filter at the top of Your Skills shows what **Needs review**.

Skills from your answers and your tools

When an answer shows a way of working you'd like to repeat, open its **…** menu and choose **Save as skill…**. Manwe puts a request in your composer; send it, and Manwe writes a skill draft for you to review.

AI tools like Claude Code and Codex can draft skills too, once you turn on **Let AI tools work on skills** in Settings › Access & approvals. Each draft says who prepared it, and stays inert until you activate it.

An agent's own procedures

Agents can develop and revise procedures during background learning. These are versioned, owned by that agent and scoped to the relevant conversation or Worlds. They do not overwrite your own skills or become available to every other agent. Inspect **Own procedures** on an agent's card. Newly written procedures are labelled untested; execution evidence is needed to show that they work.

Use it in a conversation

In chats, workflow runs and second opinions, Manwe can find and load an active skill whose description fits your task. Agents can also use applicable skills during background learning. Deep Review uses a separate research pipeline and does not currently load these skills. To pick one yourself, type `$` to see active skills for the current context and filter by name or description, or choose **Skills** from the composer's **+** menu, preview it, and press **Insert**. Use the arrow keys to choose, Return or Tab to insert, and Escape to dismiss. Nothing is sent until you send the message. `/` still opens actions and `@` opens agents and teams.

A task keeps the exact version of a skill it started with. Editing, disabling or removing the skill affects future tasks, never work already in progress. The answer's Activity shows which skills and versions were used.

Keep the instructions current

If an inactive skill fits your request, Manwe can suggest it. The review shows its source, scope and exact instructions. Choose **Enable and continue** for a skill you disabled, or **Activate and continue** for a new one. Your task can then load it and carry on. Choose **Not now** to keep it inactive and stop that recommendation for this task. Background agents can suggest a skill without waiting for a hidden review.

Select a skill to **Edit**, **Duplicate**, **Export** or switch it off. The version menu shows the active version and any draft, so you can compare before activating a change.

> **Instructions don't grant access.** A skill uses the same file, browser and connected-app permissions as the task it runs in. Scripts that come with a skill are kept for you to inspect and export; Manwe doesn't run them. Skill instructions are separate from sources and approved world memory.

STAYING ON IT

Workflows

Turn a task you repeat into a saved workflow. Keep its instructions, inputs and world together, then run it on demand or on a schedule.

Create one

Open Workflows and press **New workflow**. A workflow belongs to a world, so create one first if you haven't. Then choose how to start:

Write it myself	Fill in the editor yourself.
Browse examples	Start from a ready-made workflow for everyday life, learning, creative work, or work and community, and adjust it.
Build with an agent	Describe what you'd like to do repeatedly. Manwe asks what it needs to know and drafts the workflow for you to review.

The quickest way is often from a conversation. When an answer did something you'll want again, open its **…** menu and choose **Save as workflow**. Manwe opens a draft with your request as the instructions and the conversation's worlds and model already set.

What's in a workflow

Instructions	The result you want, written the way you'd ask in a conversation.
Home World	The world whose context the workflow uses. Add more under **Additional World context**.

Model	The workflow saves its exact model. Reasoning is captured when a run starts and reused for a valid restart. An unavailable model reports an error; changing another conversation's model does not select a substitute.
Inputs	Anything that changes between runs, like a date or a client's name. Mark an input **Required**, and give it a default for scheduled runs.
Schedule	On demand, daily or weekly, at a time and time zone you choose.

Keep methods you use across many tasks in a skill instead; any workflow can use it.

Run and review

Run now starts the workflow. If it has inputs, you see them first and press **Start run**. While it runs, the button becomes **Open run**. Every run keeps the definition and inputs it started with. If an essential detail is missing, the run asks in its conversation and continues after your reply; the header's **Activity** shows it under **Needs you**.

The list shows each workflow's schedule and status, with a filter for what's running, waiting for input, scheduled or paused. Expand a workflow and press **Open history** to see past runs, their outputs and conversations, and exactly which definition and inputs each one used. Documents and other outputs are also kept in **Settings › Vault › Saved work**.

Repeat on a schedule

Scheduled runs prepare results without you there, so they only read: they look things up and save what they find. Anything that changes something outside Manwe, like sending a message or editing a page, needs a run you start yourself, with its normal approvals. Your private replies to a run's questions never become defaults for later runs.

By default, scheduled runs happen while your Mac is awake and Manwe is open, and a missed run catches up once. To keep them going after you quit, turn on **Keep working after quit** in

Settings › General › Background work. A workflow that uses a model on your Mac still waits until Manwe is open, because the background helper doesn't load local models.

If a scheduled run can't start, Manwe tells you in a conversation and a notification. An unattended read-only run cut off by a restart can restart from its saved boundary when its definition, model, reasoning, Worlds and project context still match. Runs with saved responses, unanswered questions, approvals or uncertain effects need review; earlier actions and approvals are not replayed.

Pause schedule, in the workflow's **…** menu, stops future scheduled runs. **Stop run** stops the current one. They're separate on purpose.

> **Keep access explicit.** A workflow uses its selected worlds and the task's permissions. Saving a workflow doesn't approve future changes, or turn its output into world memory.

STAYING ON IT

Watching

Some questions don't have an answer yet. Rates move, a limit gets announced, a price drops. Ask Manwe to keep checking, and it looks things up on a schedule and tells you when there's something worth knowing.

Setting one up

Ask the way you'd ask a person. Words like "every Monday", "tell me when", "let me know if", or "keep an eye on" tell Manwe you want it to keep checking, not just answer once.

ASKS ONCE	KEEPS WATCH
What are 30-year mortgage rates right now?	Every Monday, tell me if 30-year mortgage rates drop below 6%.

Manwe answers with where things stand today, then sets up the watch. A small row under the answer shows what it's watching, how often, and when it will tell you. That first answer is the starting point every later check is compared with.

You can also start from home: pick the **Keep watch** chip under the composer (in a conversation, choose **+ › Ask differently › Keep watch** or type `/watch`) and say what Manwe should keep an eye on. If your words don't already sound like a standing request, Manwe treats it as one anyway and tells you when it changes.

> **One date isn't a watch.** "Remind me on the 15th" is a single moment, not something to keep checking. Manwe handles that kind of thing with a check-in instead. See Check-ins and outcomes.

How often it checks

If you name a time and a rhythm, Manwe uses them. If you don't, it checks at 9:00, daily for prices and rates and weekly for news and announcements. You can choose:

- **Daily** at a time you pick.
- **Weekly** on a day and time.
- **Monthly** on a day of the month. If the month is shorter, it runs on the last day.
- **Every few hours**, from every hour up to once a week.

When it tells you

A watch that reports every time it looks quickly becomes noise. So each watch has a mode that decides when a check is worth interrupting you:

Every time	Each check is reported. Right for a digest or a regular summary you want on a schedule.
When…	Only when a condition you set becomes true, like "the average rate is below 6%". You hear about it once when it turns true, and again only if it stops being true and then comes back.
When it changes	Only when something meaningful is new or different since the last check. The default when Manwe can't tell which you meant.

What each check does

Each check answers the request the way a conversation would: the same advisor, if the conversation had one, your memory from the conversation's worlds, and a fresh search for the thing being watched. The task uses its configured model; a retained advisor's fixed or allowed-model assignment is respected. An unavailable assigned model is reported instead of replaced silently.

Then Manwe compares the new answer with what the last check found. Restated facts, small wording differences, and routine noise don't count as a change. Only the comparison decides

whether you hear about it; quiet checks still update the record, so the next comparison starts from the latest facts.

Numbers are held to a stricter standard than usual. You can make it stricter still: a watch's **Evidence** setting can require a dated observation, or a current one no older than you choose. A figure in a watch update has to come from a source, cited, with its date. If no source shows the current figure, the update says so plainly instead of estimating. A watch you can't trust is worse than no watch.

The Watching place

Watching in the sidebar lists everything Manwe is checking for you, with **New watch** and a search field at the top. Each row shows the schedule and either the next check or a status: Paused, Checking…, or Needs attention when a check failed twice in a row. Expand a row to see when it reports, the last finding with the date of its evidence, and how often it has checked and reported. Its **…** menu has:

Run now	Checks right away and posts the result in the conversation, whatever the mode says.
Pause	Stops checking until you resume. Resuming doesn't make up for the runs you skipped; it picks up at the next scheduled time.
Edit	Change the name, what to check each time, when to tell you, how often, and what evidence counts. **Stop Watching** is here too. Updates already posted stay in the conversation.
Open conversation	Opens the conversation the watch belongs to, where every update lands.

You can also reach a watch from the conversation itself: the row under the original answer has a **Manage** button.

Where updates arrive

An update is posted into the watch's conversation, marked as an update, so you can reply and keep talking about it. You also get a Mac notification with a short headline; click it to open the conversation. Manwe asks for permission to send notifications the first time you set up a watch, and Settings › Notifications › **Watch updates** turns them off. Updates are saved in their conversations either way.

If you've connected Telegram, updates arrive there as well, and replying continues the same conversation. You can turn that off in Settings › Connected apps.

> **When checks happen.** Watches run while your Mac is awake and Manwe is open. Turn on **Keep working after quit** in Settings › General › Background work and they continue after you quit, as notifications from "Manwe Background". Two exceptions: a watch on a model running on your Mac waits until Manwe is open, and updates found after quit go to your Mac, not to Telegram. If the Mac slept through a check, Manwe catches up once when it's back, not once for every run it missed. If a model on your Mac is busy with something else, the check waits a few minutes for it.

STAYING ON IT

Check-ins and outcomes

Advice is only as good as how it turns out. When a decision has a date attached, Manwe comes back to it then and asks how it went. What you tell it goes on the advisor's record, so the next answer is a little better calibrated to you.

When Manwe plans a check-in

After an answer, Manwe looks at what you said for something with a date: an exam, a lease renewal, a deadline, or something you committed to doing by a certain time. If there is one, it plans a check-in for the moment it would help most. That might be two days before the exam, or the week before the lease is up. It plans at most two per answer, and most answers get none.

NO CHECK-IN	GETS A CHECK-IN
How do landlords usually decide on rent increases?	My lease renews on March 1 and the landlord wants 15% more. I'm going to try negotiating it down before then. How should I approach it?

When Manwe schedules one, a small row appears under the answer with the date and what it's about. If you don't want it, press **Cancel** on that row and Manwe won't bring it up.

After a deep review

A deep review can name a date when you'll know whether its call was right. If it does, Manwe schedules a "Looking back" check-in for that day. It reminds you what the review recommended and asks how it turned out.

When a check-in arrives

Check-ins are due at 9 in the morning on their day. Manwe looks for due check-ins every 15 minutes while it's open. If your Mac was asleep or Manwe was closed, the check-in arrives the next time Manwe is open. The background helper doesn't deliver check-ins, so this holds even with **Keep working after quit** on.

A check-in shows up in several places, so it's hard to miss:

The conversation	It's added to the end of the conversation it came from, with a couple of quick replies.
Notifications	A macOS notification with the check-in's title. Click it to open that conversation. Manwe asks for permission to notify you the first time it schedules a check-in. Settings › Notifications › **Check-ins** turns them off.
Telegram	If you've connected Telegram and left **Send check-ins and updates to Telegram** on, it's sent there too. Replying there continues the same conversation.

Before a check-in is due, you can see it coming. The home screen's **Coming up** list shows your next few scheduled check-ins, with the conversation each belongs to and when it's due. Click one to open that conversation.

Replying to a check-in

Just answer it in the conversation, the way you'd tell a friend. Make your reply the next message after the check-in, so Manwe knows what you're answering. The advisor who gave the original advice is the one who replies, since check-ins don't change who's advising you in a conversation.

If your reply says what actually happened, Manwe reads it as an outcome. Numbers and names help: "Got them down to 7% for a two-year lease" is easier to record than "it went fine". If your reply is "not now" or a new question, nothing is recorded.

How outcomes are recorded

Manwe notes the topic, what was recommended, what happened, and a lesson for next time. Then it decides how the advice held up:

Held up	Following the advice worked out.
Partly held up	Some of it worked and some didn't.
Didn't hold up	It didn't work out, or the advice was wrong.

Sometimes your reply says what happened but not whether the advice helped, or you didn't follow it. Then the outcome is noted without a verdict, and it doesn't count for or against the advisor.

A row appears under the response to your reply, saying whose record the outcome went on and how the advice held up. For a normal answer, that's the advisor who gave it. For a deep review whose panel was drawn from your advisors, it goes on each panel member's record. Otherwise it goes on Manwe's own.

Undoing an outcome

Manwe can misread a reply. Maybe you were joking, or you were describing someone else's situation. Press **Undo** on the outcome row and it's taken off the record entirely. For a deep review, that removes it from every panel member at once.

Why outcomes matter

Outcomes are how an advisor gets better at advising *you*. They feed two things:

- **The track record you see.** Each advisor's answers and card show a line like "2 of 3 outcomes held up", and the card lists recent outcomes with their lessons. See Advisors.
- **What the advisor remembers.** When an advisor answers again, the outcomes most relevant to the new question come into their thinking, along with the lessons. Advice that went badly last time shapes the next answer.

The loop only works if you close it. A check-in takes a sentence to answer. Without it, an advisor never learns whether they were right.

STAYING ON IT

Manwe on Telegram

Decisions don't wait until you're at your desk. Connect Telegram and you can ask Manwe from your phone. The answer still comes from your Mac, with your memory, your advisors, and sources you can open.

What it does

You talk to a private Telegram bot that belongs to you. Your Mac picks up each message, answers it the way Manwe would in the app, and sends the answer back. Nothing is answered by a server somewhere else, and there's no Manwe server in between: messages travel between Telegram and your Mac.

- **Ask anything** you'd ask in the app. Answers come with their sources, and the advisor who'd answer on your Mac answers here.
- **Everything shows up on your Mac.** Each Telegram conversation appears in Manwe too, marked "via Telegram", so you can pick it up later at your desk.
- **Check-ins and watch updates arrive here.** When Manwe checks in about a decision, or something it's watching is worth reporting, it lands in Telegram as well. Reply, and you continue that conversation.
- **Standing requests work too.** "Every Friday, tell me if…" sets up a watch from your phone just as it does on the Mac.

A few commands

Type / in the chat and Telegram suggests them:

/new	Starts a new conversation. Until you send it, messages continue the current one.
/second	Gets second opinions on the last question from other advisors. Missing specialists can be created when needed.

/stop	Stops the answer being written.
/help	A short reminder of what Manwe can do here.

When an answer suggests follow-up questions, they appear as buttons under it. Tap one to ask it.

Setting it up

Open Settings › Connected apps and choose **Set up Telegram**. It takes about two minutes, and Telegram bots are free. Setup walks you through three steps:

1. **Create a bot.** Open BotFather, Telegram's official helper for making bots, and send `/newbot`. Give the bot a name and a username ending in "bot". Manwe suggests both, with a button to copy each.
2. **Add its token.** BotFather replies with a token: numbers, a colon, then a long run of letters and numbers. Paste it into Manwe. You can paste BotFather's whole message; Manwe finds the token in it and checks it with Telegram right away.
3. **Say hello.** Open your bot, with the button on your Mac or by pointing your phone's camera at the code, and tap **Start**. That pairs your Telegram account with this Manwe.

Why it's private

The token works like a password for your bot, so treat it like one. Manwe keeps it in your Mac's Keychain, not in a settings file.

Pairing is what keeps the bot yours. The Start link carries a one-time code, and only the account that taps Start with that code becomes the owner. Anyone else who finds your bot and writes to it is told, at most once an hour, that this is a private Manwe. Their messages are never answered.

> **Your Mac does the work.** Manwe answers while your Mac is awake and Manwe is open. If the Mac is asleep, your messages wait in Telegram until it wakes. A Mac that's always on, like a Mac mini, answers any time. Telegram needs Manwe itself open: **Keep working after quit** doesn't cover it.

One more thing to know: if a deep review or another run is using a model on your Mac, Telegram asks you to try again when it finishes, because a local model can only do one thing at a time. Cloud models don't have that limit. See Models and providers.

Changing your mind

In Settings › Connected apps, the Telegram card shows your bot, who it's paired with, and whether it's connected.

- **Send check-ins and updates to Telegram** is on by default. Turn it off to keep those on your Mac only. You can still ask from your phone.
- **Disconnect** stops Manwe answering your bot and forgets the token. Your conversations stay in Manwe. The bot itself still exists in Telegram; to delete it, send `/deletebot` to BotFather.

If you make a new token in BotFather, the old one stops working, and Manwe tells you Telegram stopped accepting it. Choose **Set up again** and paste the new one.

SHARING

Sharing and Spaces

A decision or an answer is often meant for someone else: a partner, a client, your team. Manwe turns it into something worth sending: a styled page, a PDF, a slide deck, or a link. A Space collects a world's best work into one site.

Share an answer or a decision

Open an answer's **…** menu and choose **Share & export…**. For a deep review, use the export menu in the decision document's header. The studio opens with a live preview on one side and the look on the other.

Template	**Editorial**: serif headlines and magazine rhythm. **Swiss**: a strict grid, numbers forward. **Executive brief**: an answer-first memo for clients and partners. **Studio**: a bold cover, large type and cards.
Accent	A swatch, your own color code, or the template's.
Fonts	Five free pairings, embedded in what you export so it looks the same everywhere.
Appearance	Auto follows the reader's light or dark setting; or fix it to Light or Dark.
Byline	Your name or your company's, shown at the top and bottom.

Your choices are remembered, per world, for next time. Then pick what to make:

PDF	Real pages with selectable text and working links.

HTML	One self-contained page, fonts included, that opens in any browser.
Slides	**PowerPoint**, **Keynote**, or a widescreen **Slides PDF**. Keynote export needs Keynote installed, and asks once for permission to control it.
Share…	The Mac's share menu, for Mail, Messages, AirDrop and the rest, with the PDF and the link if there is one.
Get link	A web address anyone can open. See below.

Copy Markdown, in the same menus, copies the plain text.

Links

A link is served from your Mac, through Cloudflare, to whoever opens it. Nothing is uploaded to Manwe, and there's no account to make. Press **Get link**; once Manwe confirms the link works, **Copy link** appears.

A link is a snapshot. When you change the look or the content, press **Update** to publish it to the same address. **Turn off** stops the link working for everyone, and deletes the published copy.

Links are long and random, so they can't be guessed, and they ask search engines not to list them. Still, anyone you send one to can pass it on, so share accordingly.

Keeping links working

Links open while your Mac is awake and sharing is on. Manage it all in Settings › Sharing:

Let people open your links	The main switch. Off, your links stop opening until you turn it back on.
Keep links working after you quit Manwe	Hands the links to Manwe's background helper, the same one that runs scheduled workflows. Without it, links only work while Manwe is open.

Shared now	Everything you've shared, with **Open**, **Copy** and **Turn off**.

> **Why the address can change.** Without any setup, links use a free Cloudflare address that's new each time sharing restarts: when your Mac restarts, or when Manwe restarts if the helper is off. Your pages stay shared; copy the link again to get the new address.

Your own web address

For links that never change, use a domain you own, like `share.yourname.com`. You need a free Cloudflare account with that domain in it. Under **Your own web address**, type the address and press **Set up with Cloudflare**. Cloudflare opens in your browser; sign in and pick the domain, and Manwe creates the connection and points the address at your Mac. Already made a tunnel in Cloudflare yourself? **Use a tunnel you made in Cloudflare** takes its token instead.

Spaces

Every world, and every Moon, has a **Space**: a small site that brings its decisions and answers together, for a client, a partner, or your future self. Open the world in Worlds and choose the **Space** tab.

01 **Fill it**
Under **Decisions in this World**, press **Add** beside finished deep reviews. Add answers from a conversation with **Add to Space** in their **…** menu.

02 **Shape it**
Give it a title and a short introduction for the cover, and put the items in order.

03 **Style it**

Open **Look** for the same templates, accent, fonts and byline as single pages.

---

04 **Share it**

Get link publishes the site; puts the cover and every item in one **PDF** document.

The published Space opens on a contents page with a card for each item. Each item has its own page, with Previous and Next to read through in order.

The preview always shows the items as they are now. The link changes only when you press **Update**, so you decide when the people you shared it with see something new.

SETUP

Models and providers

Manwe doesn't come with one fixed brain. It can run a model on your Mac, use an AI subscription you already pay for, or call a provider with your own API key. Which you choose shapes privacy, speed, and how much an answer can do.

Three ways to run Manwe

On this Mac	A model downloaded to your Mac and run on Apple Silicon. Private and free to use, but slower, and it handles one thing at a time.
Your subscriptions	Claude and OpenAI models through the command-line tools you sign in to: Claude Code and Codex, using the plan you already have.
API keys	Direct access with a key from Anthropic, OpenAI, Alibaba, or another compatible provider. Experimental.

Everything is in Settings › Models & providers, under **Intelligence**. The **Images** tab, covered below, is separate.

Models on your Mac

Download one under **On this Mac**. Each model lists what it's good at and how much memory your Mac needs:

Everyday	Gemma 4 E4B. The smallest curated download for everyday questions. Needs 8GB of memory or more.

Balanced	Qwen3.5 9B or Gemma 4 12B. Larger local options; speed and task quality depend on the model and hardware. Needs 16GB or more.
Large	Qwen3.8 27B. A larger option for demanding tasks; accuracy and speed depend on the task and hardware. Needs 32GB or more.

Downloads are several gigabytes each. You can delete a model and download it again later. Models you downloaded before this list changed stay usable under **Earlier downloads**.

Models you already have

If you already run models on your Mac, **Your installed models** brings them in without downloading again:

Choose model folder…	Links an MLX model folder where it is. **Remove link** forgets it; your files are never deleted.
Connect local server…	Uses a model served by LM Studio, Ollama, Prism or another OpenAI-compatible runtime on your Mac, including GGUF models. Start the server yourself; Manwe doesn't launch it.

The point of a local model is that the model itself never leaves your Mac: your question, your memory, and the answer are all worked out on your own hardware. When a question needs current facts, Manwe still searches the web for them, so the search goes out even though the thinking stays home.

Subscriptions you already have

If you use Claude Code or Codex, Manwe can use your subscription. Open Settings › Models & providers. Each tool shows whether it's installed and signed in, and lists the models it reports. **Refresh** picks up new ones as the tool adds them.

- **Claude Code** gives you Claude models through your Claude subscription.
- **Codex** gives you OpenAI models through ChatGPT or your OpenAI account.

- **Qwen Code** may retain saved model choices, but its current task execution route is not verified. Manwe reports that limitation rather than running an older route or substituting another provider.

Manwe runs these tools on your Mac and relies on their own sign-in. If a tool isn't signed in, **Sign in** opens Terminal with that tool's login, and you sign in there. Manwe never sees your password, and it doesn't hold a separate account with any of these companies.

API keys

Further down Settings › Models & providers, you can paste an API key for Anthropic, OpenAI, or Alibaba's DashScope. **Add provider** adds as many more as you need: one click for OpenAI, Anthropic, Mistral, DeepSeek, Kimi, Z.ai, OpenRouter, Groq, Together AI, Fireworks AI, Ollama or LM Studio, or **Custom endpoint…** for any service that speaks the OpenAI format. **Fetch models** lists what each one offers, and you can pick its default model.

Keys are kept in your Mac's Keychain. API connections are experimental. A listed model or successful connection does not prove that a model can complete a particular task. Tool-call formatting, reasoning controls and output quality depend on the exact model and provider configuration.

Choosing a model

The model menu sits in the composer, next to where you type. It lists every model that's ready to use: downloaded, signed in, or set up with a key. Search it with **Find a model**. Pick one and Manwe uses it for what you ask next, and keeps using it in that conversation until you pick another. New conversations and watches start with the **Default model** from Settings › General.

Models for individual agents

An agent can follow the conversation model, use one assigned model, or use an allowed list, each with its own reasoning setting. Configure this in **Agent settings** on its card. Groups and review participants keep those assignments; an unavailable assigned connection reports an error rather than silently switching. See Advisors.

Reasoning effort

Many models can think longer before they answer. Choose how much with **Effort** in the composer's model menu, or **Reasoning** in Settings. The levels depend on the model: cloud models offer steps like Low, Medium and High, and models on your Mac can switch thinking on or off. Default uses the backend's default; for the built-in local models, thinking is off by default. Higher is slower and often sharper. It's set per model, so you can keep a fast model quick and let a careful one take its time.

Model size and prompt checks do not guarantee answer quality. Local model support is still being refined; a model can fail during preparation or generation. Check the message's activity for the actual error, and choose another model explicitly if you want to continue elsewhere.

Images

Manwe can make images when you ask for one. Choose providers on the **Images** tab of Settings › Models & providers: your Codex subscription, Apple Image Playground on Macs with Apple Intelligence, or an API key for OpenAI, Google Gemini, Black Forest Labs FLUX, Ideogram, Recraft or Stability AI. When an answer needs an image, Manwe asks which provider to use; your conversation's model stays the same, and every paid generation asks for approval. Images land in the answer and in Saved work, with **Save image…**.

How the choice shows up

A model on your Mac can only produce one thing at a time. A cloud model can handle several requests at once, and that changes how some features feel:

- **Second opinions and deep reviews.** On a cloud model, advisors answer at the same time, and a deep review's panel speaks at once within each stage. On your Mac they take turns, so it takes longer.
- **Waiting your turn.** While a run is using a local model, a watch waits for it, and a question from Telegram is asked to try again later.
- **After you quit.** The background helper that keeps workflows and watches going after you quit uses cloud models only. Anything set to a local model waits until Manwe is open.
- **What an answer can include.** Answers can carry tables, charts, calculators and explorers, and Word, Excel and PowerPoint files you can preview and download, whatever the model.

Larger models tend to build them better.

> **A rule of thumb.** Use a local model when privacy matters most or you'd rather not use a subscription. Use a cloud model for decisions where depth and speed matter, especially deep reviews.

SETUP

Connecting apps

What you know is spread across your Mac: meeting notes in one app, files in another, a long conversation with an AI coding tool in a third. Connections let that knowledge flow in both directions, so Manwe can look things up in your apps, and your other AI tools can stop asking you to explain yourself again.

Set everything up in Settings › Connected apps: online services, Apple apps, apps on this Mac, and Telegram. Control what outside AI tools can read or ask Manwe to do in Settings › Access & approvals.

Connected services

Notion, Linear and Slack each have a row. For Notion and Linear, press **Connect** and sign in through your browser. Slack needs a one-time **Set up** first, because each workspace registers its own Slack app; the setup sheet and its guide walk you through it.

Once connected and assigned to the selected context, Manwe can search and read them when a question needs it, and prepare changes: new or updated Notion pages, Linear issues and comments, Slack messages and canvases. The row shows how many tools the account offers, and **Permissions and authorization** shows what it can reach. **Disconnect** stops new use.

Recent app changes lists what Manwe changed and whether each change was confirmed. When Manwe couldn't confirm one, press **Review result**, check the app, and say whether you found the change.

Choose where a source is available

Under each connected account or local MCP app, use **Available in** to choose its Worlds. Select **Chats** explicitly if it should also be usable in conversations outside any World. A connection

with no assignment stays connected but is unavailable to chats. A work account assigned to one World does not supply context to an unrelated World.

Use parent sources separately lets a Moon read its parent's assigned sources. That access is read-only; making changes needs a direct assignment to the Moon. Parent source access is independent of **Use parent context** for memory. Clearing an assignment or disconnecting stops further use in that scope.

Apple apps

Calendar, Reminders, Contacts, Notes, Mail, Maps, Photos, Music and Shortcuts work without setup; macOS asks for permission the first time. Each has a switch in the **Apple apps** card. See Apple apps for what each one can do.

Apps on this Mac

Many apps and tools share their data with AI tools through MCP. When one is connected, Manwe can look things up in it when a question needs your own meetings, notes, or files, the same way it searches the web when a question needs current facts.

How Manwe finds them

Manwe lists every MCP server set up in Claude Code, Claude Desktop and Codex, plus apps in your Applications folder that come with an MCP helper built in. When you install a new one, Manwe asks once whether to connect it. Apps that were already there when you started are simply listed, not prompted.

Some entries carry a note in plain words, so you know what connecting means. A server might download code when it starts, run a command that isn't part of an app, or pass settings from your AI tool's configuration, which may include a key. It's your choice: connect the ones you trust.

Connect and disconnect

Each app has a **Connect** button. Connecting starts the app's helper and reads the tools it offers. The row then shows how many of them read, and how many make changes. **Disconnect** stops Manwe using the app at any time.

Changes follow your approval mode

Reading is direct. A change, in a connected service, an Apple app or an app on your Mac, follows your **Action approvals** setting. Under **Ask for approval**, the card explains the action in plain language, with exact details available to expand. **Approve**, **Deny** and, for eligible actions, **Always allow similar** control it. **Approve for me** authorizes requested routine changes within the task's access. **Full access** permits work within granted access, too; destructive changes still require explicit approval in every mode. See Tasks and approvals.

A change always needs a conversation where you can see it. Scheduled workflows and watches only read, and if you deny a change, Manwe doesn't try it again.

AI tools reading Manwe

The other direction, configured in Settings › Access & approvals. Manwe ships a small MCP helper that lets AI tools on your Mac read what Manwe knows about you: your profile, your worlds, your confirmed memory, and past decisions. Ask Claude Code for help with a project, and it can check your goals and constraints instead of asking. With your permission, it can also hand Manwe work to do.

Under **Tools on this Mac**, Manwe lists Claude Code, Codex, and Claude Desktop, with whether each is installed and connected. **Connect** registers Manwe with that tool for you; **Disconnect** removes it; **Reconnect** fixes a registration that points somewhere else.

What they can see and do

Everything here is off until you turn on **Let other AI tools read Manwe** (connecting a tool turns it on). Then each part is its own switch:

Let other AI tools read Manwe	Your profile, worlds, confirmed memory, and past decisions. Drafts and suggestions are never shared, because you haven't approved them.
Share past decisions	On by default. Decision records are labelled as Manwe's analysis, not as facts you confirmed.

Let AI tools delegate tasks	Off by default. Lets a tool ask Manwe a question, research, convene a panel, or hand over a whole task that uses Manwe's browser, connected services and saved work. Tasks run in Manwe with your model and your approval settings; questions and approvals appear in Manwe, never in the other tool.
Let AI tools file suggestions	Off by default. What a tool learns about you lands in memory review as a suggestion, and it can prepare a whole world for review. Nothing becomes memory or a world until you approve it. See Memory.
Let AI tools work on skills	Off by default. A tool can read your skills and add or edit them as drafts. A draft does nothing until you activate it. See Skills.

> **Why suggestions, not memory.** Another tool can misread you, or remember something that was true for one afternoon. Making its findings wait for your approval keeps your memory yours: Manwe only acts on what you've confirmed.

Other MCP tools

Using an AI tool that isn't on the list? Choose **Copy setup JSON** under **Other MCP tools**, and paste it into that tool's MCP settings. It points the tool at Manwe's helper on your Mac. The same switches above decide what it can see and do.

SETUP

Apple apps

Your calendar, reminders, notes and mail already know a lot about your week. Manwe can work with the apps on your Mac directly, so "what's on Thursday?" or "add a reminder to call the bank" just works.

Turning them on

There's nothing to install. The **Apple apps** card in Settings › Connected apps lists each app with a switch, on by default, and its status:

Asks on first use	macOS will ask for permission the first time Manwe needs the app.
Allowed	You've given permission.
Not allowed in System Settings	You said no earlier. **Open Settings** takes you to the right place in Privacy & Security to change it.
Ready	Needs no permission.

Switch an app off and Manwe stops using it right away.

What each app can do

Calendar	Read your schedule for any range up to a year, and add events.
Reminders	Read your lists, add reminders with a due date, and mark them done.
Contacts	Look people up by name or email. Read only.

Notes	Search and read your notes, and create new ones.
Mail	Search and read your inbox, and prepare a draft that opens in Mail for you to send. Manwe never sends mail.
Maps	Find places, and travel times by car, on foot or by transit.
Photos	Find photos by date or album. Manwe sees the dates, places and names, not the pictures.
Music	See what's playing, and play, pause or skip.
Shortcuts	List your shortcuts and run one by name. A shortcut can do anything it was built to do; running it follows your approval mode.

Changes follow your approval mode

Reading happens directly. Adding an event, creating a note, completing a reminder, opening a mail draft, controlling Music or running a shortcut follows your **Action approvals** setting. **Ask for approval** shows the action first. **Approve for me** allows requested routine changes within granted access; **Full access** also acts within that access. Destructive actions still require explicit approval. See Tasks and approvals. If you deny it, nothing changes and Manwe doesn't try again.

If a change may have only partly worked, Manwe says so instead of guessing, so you can check the app before trying again.

> **Only while Manwe is open.** Apple apps work in your conversations. Scheduled workflows and watches running after you quit don't use them. What Manwe reads from an app goes to the model you're using, like any other context; place searches go to Apple Maps.

SETUP

Your data

Manwe gets more useful the more it knows about you, which means it ends up holding things that matter. This page covers where that lives, what leaves your Mac and when, and how to make sure you never lose it.

Where it lives

On your Mac. Your conversations, memory, worlds, advisors, decisions, saved work, Spaces, the pages you've shared, and the documents you've attached are kept in Manwe's own folder on this Mac, along with any models and tools you've downloaded. There's no Manwe account to sign in to, and your record isn't copied to a Manwe server.

Secrets are kept apart. API keys, sign-ins for Notion, Linear and Slack, your Telegram bot's token and your Cloudflare tunnel token are stored in your Mac's Keychain, not alongside the rest. Settings that apps on your Mac use, like keys in Claude Code's configuration, are read from there when needed and never copied into Manwe.

What Manwe learns about how you work stays here too: the corrections and choices it learns from, and the small models it trains on them, are kept in Manwe's database on this Mac and included in your backups. They are never sent to a model provider. The preferences learned from them become memory, and relevant memory is part of the context sent to the model you choose. Turn learning off or reset it in Settings › Profile › Memory and learning; see Memory.

Manwe keeps its own conversation history. Its Codex exchanges request temporary provider threads and reconstruct the relevant Manwe history for later replies; they do not reuse a saved Codex thread ID. Existing history in Codex is not deleted.

What leaves your Mac

That depends on what you've set up, and it's worth knowing exactly:

Cloud models	When you use a model through Claude Code, Codex, or an API key, the provider receives what the model needs to answer: your question, the profile and memory relevant to it, sources Manwe found, document text and supported images supplied for the task. Each participating agent uses its assigned connection, so a mixed-provider group may send relevant task context to more than one provider.
Research	When a question needs current facts, Manwe searches for them, so its search queries go out to the research sources it uses.
Connected services	Notion, Linear and Slack receive Manwe's searches and reads, and any change you approve.
Apps on this Mac	Requests go to the app's helper on your Mac. Some helpers reach the internet themselves; their notes in Settings say so.
Apple apps	Calendar, Notes, Mail and the rest stay on your Mac, apart from what goes to your model as context. Place searches and travel times go to Apple Maps.
Shared links	Pages you share are served from your Mac through Cloudflare to whoever opens the link. Nothing is uploaded to Manwe. See Sharing and Spaces.
Telegram	If you've connected it, your messages and Manwe's answers pass through Telegram on their way between your phone and your Mac.
iCloud Drive	Only if you turn on iCloud backups, below.

Usage analytics	On by default, and you can turn it off at any time in Settings › Access & approvals › Usage data. Manwe sends TelemetryDeck a small, fixed set of events: the app opening, onboarding finishing, which screen you open, a conversation starting, and whether a task completed, failed or stopped. Each comes with the app and macOS version and a random identifier that isn't linked to you. Never your chats, documents, file paths, names, credentials, memory or model responses. Turning it off stops sending immediately and deletes the random identifier. Like any web request, TelemetryDeck sees your IP address while an event is delivered.
Apple Maps	When an answer shows a map, Manwe asks Apple Maps to find its places and work out its driving and walking routes. Apple receives the place names and the route start and end points, not your conversation. Apple’s Maps privacy terms apply to those requests.
Voice prompts	Nothing. When you press the microphone in the composer, Manwe listens until you press it again and turns your speech into text with Apple’s on-device speech recognition. Your voice never leaves this Mac. Pick the language in Settings › General › Voice prompts. The first time, macOS may download Apple’s speech model for that language. If your Mac can’t recognise your language on device, voice prompts stay off rather than using a server.
Downloads and updates	Downloading a model fetches it from Hugging Face. Manwe checks for app updates about twice a day, and while sharing is on, its sharing tool checks for its own updates at most once a day. These requests carry nothing from your record.

With a model running on your Mac, the model itself never sends anything anywhere. See Models and providers. To share Manwe's knowledge with other AI tools on your Mac, see Connecting apps; that's off until you turn it on.

Enabled agent background learning can also use configured providers and public research while Manwe is open and online, without a new message from you. Its daily allowance and

pause controls are separate from requested work.

Backups

Manwe's own folder is the kind of place an uninstaller cleans out. So Manwe keeps copies of your record somewhere else, where deleting or reinstalling the app can't reach them. Find them in Settings › Data & backups.

- **Automatic backups** run every 6 hours when something changed, into **Documents › Manwe Backups**. If nothing changed, no new copy is made.
- **Back up now** makes one whenever you like, for example before a big change. **Show in Finder** opens the folder.
- **Documents you attached**, the Vault and your saved work are backed up too, not just the database. Spaces and shared pages are included. Temporary tunnel state is excluded; restoring leaves sharing off until you turn it on again.

Manwe thins out old automatic backups so the folder doesn't grow forever. It keeps every backup from the last two days, one a day for a month, and one a month for a year. Backups you made yourself are never thinned, and neither is the newest backup or the one holding the most.

iCloud Drive

Also back up to iCloud Drive puts copies in iCloud Drive › Manwe Backups as well, so they survive losing the Mac. This uploads everything, including all your memory, so turn it on only if you're comfortable with that. It's off by default, and unavailable if iCloud Drive is off on this Mac.

Restoring

Recent backups are listed with their date and what they hold: how many decisions, worlds, and memory items. Choose **Restore…** on one, and Manwe quits and reopens with that backup in place. Your current database isn't deleted; it's moved to Manwe Backups › Replaced Databases, so you can go back. A run in progress stops. Restoring is postponed while **Keep working after quit** is on; Manwe tells you to turn it off first. Background tasks come back as saved progress, and you continue them yourself.

A backup stored in iCloud that isn't on this Mac yet shows **Download** first. And if Manwe ever opens empty, after a reinstall or an accident, it looks for a backup and offers to restore it.

> **Backups belong to you.** They're ordinary files in your Documents folder. Copy them to a drive, or include them in Time Machine, the way you would anything else you can't afford to lose.

Vault

When you attach a document to a question, like a PDF, Word, Excel or PowerPoint file, a CSV, or text, Manwe keeps a copy in the Vault, in Settings › Vault. You can also drop files straight into the Vault. Its **Saved work** tab keeps the documents, spreadsheets and other outputs Manwe made for you, with their versions. Manwe pulls out the text and splits long documents into sections, so it can find the passage a question needs instead of rereading the whole thing.

Anything in the Vault can be attached again later without finding the file. Removing an item takes it out of the Vault and the attach menus; memory you already approved from it stays.

Archive

If you used Manwe before 1.0, your earlier runs are in Settings › Archive. They're read-only: you can search them and open any one to read its report and how the panel got there. Nothing new is made this way, but Worlds still learns from their predictions. Deleting a run removes it from this Mac; backups keep their copies.

Privacy, in short

- Your record lives on your Mac, with copies where you choose.
- Your memory goes to whichever model you use, on this Mac or in the cloud. Cloud providers see it only when you use them.
- Proposed facts about you do not become approved personal or World memory without review. Agents keep separate scoped learning records. See Memory.
- Other AI tools see your confirmed memory only if you let them read Manwe.

- App changes follow the selected approval mode. Approve for me permits requested routine changes; destructive actions require explicit approval in every mode. Mail only prepares drafts.
- A shared link opens for anyone who has it, until you turn it off.
- Basic usage analytics are on by default and never include your content. Turn them off in Settings › Access & approvals.

SETUP

Make it yours

You'll spend a lot of time reading in Manwe, so it should read the way you like. Settings › Appearance changes the look, the type and the window, and every change shows at once.

Look

Mode	**Match Mac** follows your Mac's light or dark setting; **Light** and **Dark** stay put.
Theme	Eight complete looks, each with a light and a dark version: Manwe, Graphite, Paper, Varda, Taniquetil, Forest, Ember and Riso. Each card shows a small preview of the window.
Accent	Keep the theme's own colour or pick another. Buttons, selections and the first colour in charts follow it, so charts and interactive answers belong to your theme too.
Window transparency	Lets your desktop show softly through the window. The slider sets how much. The composer becomes frosted glass and the conversation fades out above it, so text stays readable. If **Reduce transparency** is on in your Mac's Accessibility settings, the window stays solid.

Type

Interface	The font for menus, the sidebar and controls.

Reading	The font for answers. It can follow the interface font or be any font installed on your Mac; each one is listed in its own lettering so you can see it before you pick.
Code	The font for code and other fixed-width text.
Text size	Four steps, from compact to large, applied across the app.

The preview under the fonts shows an answer in your choices, with a line of code.

Advisors

Printed texture gives advisor glyphs a fine, printed grain. Turn it off for flat colour.

> **Nothing to lose. Reset to default** at the bottom puts everything back to the Manwe theme. Your choices never change your conversations, memory or saved work.