

USER MANUAL

The Complete Guide to *Sal*.

Everything you need to know about your private, local AI assistant for Mac —
from first launch to advanced features.

Sal v2026.26.51 · June 2026

Activation

When you purchase Sal from hellopartner.ai, you'll receive an email from LemonSqueezy containing your personal **license key** and a secure download link.

Steps to Activate

1. Download the `.dmg` file from your purchase email.
2. Open the disk image and drag **Sal** into your **Applications** folder.
3. Launch Sal from Applications (or Spotlight).
4. On the activation screen, paste your license key from the LemonSqueezy email.
5. Sal will validate your key online and activate instantly.

Tip: Your license key is tied to your machine. If you need to move Sal to a new Mac, see [License Transfer](#).

System Requirements

COMPONENT	REQUIREMENT
Processor	Apple M3 or later
Unified Memory	16 GB minimum (24 GB+ recommended)
macOS	The latest version of macOS — currently macOS Tahoe 26.4.1
Disk Space	~22 GB for AI models (one-time download)
Internet	Required for activation and model download; optional thereafter

Why Apple Silicon only? Sal runs AI models directly on your Mac's Neural Engine and GPU using Apple's MLX framework. Intel Macs lack the unified memory architecture required for efficient local inference. We're currently testing on M3 — support for M1 and M2 may be added in a future update.

You must be on the latest macOS. Sal requires the current release of macOS — **macOS Tahoe 26.4.1** as of this writing. Sal tracks the latest Apple release and will not run on older versions. Before installing, go to **System Settings → General → Software Update** and make sure your Mac is fully up to date.

Model Download

On first launch after activation, Sal downloads the core AI models she needs to run locally. This is a one-time process — approximately ~5 GB of model weights.

What Happens

- Sal displays a progress screen showing each model as it downloads.
- Downloads resume automatically if interrupted.
- Models are stored locally in Sal's application support folder.
- Once complete, Sal warms the models into memory and becomes ready.

How long? The core download typically takes just a few minutes. After this, Sal is ready to chat, search the web, manage your files, calendar, and reminders — all offline. The larger image generation model (~24 GB) begins downloading automatically in the background once the core models are ready. You can start using Sal immediately — you don't need to wait for it.

Permissions

After downloading her AI models, Sal will immediately launch the **Privacy & Permissions** screen. There are a number of permissions requested — and we understand that can feel like a lot. However, this access is precisely what makes Sal so remarkably useful.

To ensure a smooth, uninterrupted experience, we highly recommend opting in to all settings right away. Because Sal runs entirely locally, your privacy is guaranteed: **none of your data is shared or ever leaves your device.**

Streamlined Permission Management

We've eliminated the confusing cascade of macOS dialogs, replacing them with a fully transparent system:

- **The Privacy Screen:** Lets you easily review and approve all required permissions at once during initial setup.
- **Settings Panel Control:** You can review, grant, or revoke your permissions at any time. Just open the main **Settings** window, navigate to the **Privacy & Access** section, and click **Manage Access**.
- **Contextual In-App Alerts:** If you ask Sal to perform an action but haven't granted the necessary permission, she will explicitly tell you. A simple in-app alert will appear, allowing you to grant the required permission right there in context.

Additionally, macOS will occasionally prompt you natively for access to your Documents, Desktop, or Downloads folders as Sal indexes your files in the background. These localized alerts are completely safe and normal.

Navigating macOS System Settings

macOS has strict security protocols, and sometimes you may need to manually grant permissions in **System Settings** → **Privacy & Security**. If Sal isn't showing up in the list for a specific entitlement (like **Full Disk Access** or **Accessibility**), you will need to add her manually:

1. Open macOS **System Settings** and navigate to **Privacy & Security**.
2. Select the specific entitlement you need to grant (e.g., **Full Disk Access**).
3. Click the small **plus (+) symbol** at the bottom of the application list. (You may need to authenticate with your Mac password or Touch ID).
4. In the file browser that appears, navigate to your **Applications** folder, select the **Hello Partner** app, and click Open.
5. Ensure the toggle switch next to the app is turned **on**.

Why So Many?

For Sal to be truly helpful — not just a chatbot, but a real assistant — she needs to interact with your actual life: your files, your calendar, your reminders, your screen. Every permission you grant unlocks a real capability.

PERMISSION	WHY SAL NEEDS IT
Screen Recording	Ephemeral Sight — so Sal can glance at your screen when you ask
Microphone	Push to Talk — so Sal can hear you when you press the button

PERMISSION	WHY SAL NEEDS IT
Camera	Visual input — so Sal can see what you show her
Calendar	Schedule awareness — so Sal can read and manage your events
Reminders	Task management — so Sal can create and check your reminders
Full Disk Access	File operations — so Sal can find, read, create, and organise files across your system
Accessibility	Window management and system interaction

Screen Capture Setup: It is a 2-step process. To enable screen capture, you'll need to grant permission in **System Settings → Privacy & Security → Screen Recording**. macOS requires a one-time restart after enabling.

Why It's Safe

Here is the difference between giving these permissions to Sal versus giving them to any other app:

Sal lives on your computer. Not in the cloud. Not on someone else's server. When she accesses your calendar, your files, your screen — that information stays on your machine. Period.


HelloPartner.ai doesn't even have a server that could receive your information. There is no backend, no analytics endpoint, no telemetry pipeline. We couldn't collect your data even if we wanted to — and we don't want to. This is the core of our mission. It is what makes us different.

Our promise: Whatever Sal sees on your machine stays on your machine. She does not send any data anywhere. Ever. To get the most out of Sal, we highly recommend granting all requested permissions — you're trusting your own computer, not a company's server.

Opening Sal

Sal lives in your **menu bar** — she's always one keypress away.

Keyboard Shortcut

Press   (Option + Space) at any time to toggle Sal's panel. This works from any app, any fullscreen space, any context.

Menu Bar Icon

You can also click Sal's orb icon in the macOS menu bar to open or close the panel.

Always available. Sal is designed to be a presence, not just an app. She's there when you need her, invisible when you don't.

Background Notifications

If you close Sal's panel while she's still thinking, you don't need to keep checking back. Sal will let you know when her response is ready using two indicators:

- **Green dot on the menu bar orb:** A small green dot appears on Sal's orb icon whenever a response — or a Goal Seeking job — completes while the panel is closed. The dot is always visible and clears automatically the moment you open Sal.
- **macOS notification:** A system notification is delivered to Notification Center with the title "*Sal has responded*" and a preview of the first line. For Goal Seeking completions, the notification reads "*Goal Seeking Complete*" with the phase count.

No setup required. Sal will request notification permission on first launch. If you decline, the green dot still works independently — it's the primary, always-reliable indicator. You can manage notification preferences at any time in **System Settings → Notifications → HelloPartner.**

Chat

The heart of Sal is a natural-language conversation. Type anything into the input bar and press `Return`. Sal processes your message entirely on your Mac using local language models — no cloud, no API calls, no data leaving your machine.

What You Can Do in Chat

- Ask questions on any topic
- Get help writing, editing, and summarizing text
- Brainstorm ideas and get feedback
- Manage your files, calendar, and reminders through natural language
- Generate images and search the web

Tip: If Sal is in the middle of a response and you want to stop her, just send a new message. That interrupts her immediately.

Message Menu

Every message — yours and Sal's — has a ⋯ menu. Click it to:

- **Copy Markdown:** Copy the message content as formatted Markdown.
- **Speak:** Have Sal read that specific message aloud.

Sal maintains conversational context within each session and draws on her long-term memory across sessions.

Thinking Mode

Sal can think deeply about complex problems. Unlike other tools, there are no complicated settings to turn on. All you have to do is ask her to think.

How It Works

Just ask Sal to think when you send your message. You can say something like, *"Think step-by-step to solve this problem,"* or *"Can you think about how to structure this?"*

Sal will understand what you want and adjust how she works. A thinking bubble will appear, and you can watch as she works through the problem. It is a fascinating process to observe.

Speed When You Need It

For those times when you don't need Sal to think deeply, just ask your question normally. She will give you great, fast responses without the extra time spent on thinking. It is an elegant design that gives you deep thought when you need it, and immediate speed when you don't.

Conversations & Memory

Unlike cloud AI services with enormous context windows, Sal runs locally on your hardware — which means she keeps her attention tightly focused on your **current conversation** rather than trying to hold everything in memory at once. This is what makes her fast and responsive.

Short-Term vs. Long-Term Memory

What you see in Sal's chat window is her **short-term memory** — the messages she's actively thinking about right now. Sal manages this herself, adjusting how much she holds based on what your Mac can handle.

You can scroll back through your full message history directly in Sal's response panel — there's no separate window to open. As you scroll, your original message for each turn stays pinned at the top so you always know what you asked. Every conversation is safely stored in Sal's **long-term memory**, and nothing is ever lost.

Long-Term Recall

If you need to revisit something from a past conversation, just ask:

- *"Do you remember what we talked about regarding the Q3 budget?"*
- *"What did I ask you about last Tuesday?"*
- *"Try to recall our conversation about the kitchen renovation."*

Sal will search her long-term memory by topic or timeframe and bring back what she finds. She won't always get it perfectly — but she gets better the more specific you are.

Session Commands

Most of the time, you won't need these — Sal manages her own memory just fine. But they're there when you want more control.

- `/new session` — Clears Sal's short-term memory and starts fresh. Use this when you're changing topics and need Sal's full attention on something new — especially after a long session. You're essentially saying: *"Forget everything we've been talking about. Focus on this."*
- `/restore session` — Brings back as much of your most recent conversation as Sal can fit. Useful after a `/new session` when you want to pick back up where you were. Keep in mind: if your new session was long, the restore may not be able to bring back much of the earlier conversation.

These two commands often go hand in hand. Start a new session to focus, do your work, then restore when you want to get back to what you were doing before.

Topics Explorer

As you chat, Sal keeps session notes and tracks topics to manage her memory behind the scenes. When your conversation grows too long to fit into her active memory budget, Sal automatically groups the oldest parts of the chat into summarized units called **topics** to keep the conversation fast and responsive.

This is a key transparency feature: you can actively see and read exactly what Sal has recorded in her session notes, giving you complete visibility into her understanding and memory. You don't ever need to look at or use this feature—Sal handles memory eviction and retrieval seamlessly on her own.

However, if you want to hyper-focus Sal on a specific historical subject, or explicitly control what she remembers right now, the **Topics Explorer** gives you manual control over her memory state.

How to Open

You can toggle the Topics Explorer card in the Sal interface in two ways:

- **From the Menu Bar Orb:** Click the Sal orb in your macOS menu bar, hover over **Session**, and select **Toggle Topic Explorer** (shortcut:  when the menu is open).
- **From the Main App Menu / Keyboard Shortcut:** When Sal's window is active, select **Session** → **Toggle Topic Explorer** from the macOS menu bar at the top of the screen, or press  (Cmd + T).

Active vs. Past Topics

Topics exist in one of two states inside the explorer:

- **Active Topics:** These are currently loaded in Sal's memory. She has passive awareness of these topics and can recall them instantly during your conversation.
- **Past Topics:** These are archived segments of your conversation history. Sal doesn't actively remember them right now, but they are stored locally on your device. You can manually re-activate them, or Sal can retrieve them programmatically if you ask her to.

Using the Topics Explorer Card

The Topics Explorer appears as a floating card in the Sal interface that you can expand or collapse:

- **Collapsed State:** A minimal header bar showing the active topic count (e.g.,  **Active Topics**) with a chevron to expand and a close button () to hide it.
- **Active Topics Zone (Top):** Displays the list of currently active topics. Click the minus button () next to any topic to deactivate it and move it to Past Topics.
- **Past Topics Zone (Bottom):** Lists all archived topics not currently in Sal's active memory. Click the plus button () next to any topic to re-activate it and inject it back into her memory.
- **Zone Resizer:** Drag the horizontal divider bar between the Active and Past zones up or down to adjust their heights. Sal remembers your size preferences.

Searching and Sorting Past Topics

To help you locate historical context in longer sessions, the Past Topics zone includes:

- **Search Bar:** Filters past topics in real-time. This is a semantic search that matches against topic titles, summaries, keywords, and core key facts.
- **Sorting Menu:** Sorts past topics by **Recent** (newest first), **Chronological** (oldest first), or **Alphabetical** (by title).

Topic Deep Dive (Accordion View)

Click on any topic title inside the explorer to expand it and view details generated by the AI during the archiving process:



- **Summary:** A single-sentence gist of what was discussed and the outcomes reached.
- **Keywords:** Key terms, entities, and technical concepts mentioned.
- **Key Facts:** A bulleted list of critical facts, constraints, decisions, or user preferences established.
- **Source Messages:** Click this accordion to expand and read the actual chat transcript for that topic (cleaned of raw code/JSON logs for easier reading). Clicking the ellipsis button (`...`) next to any message lets you:
 - **Copy Message:** Copy the text to your clipboard.
 - **Save Message:** Export the message as a Markdown (`.md`) file on your Mac.
 - **Speak Message:** Listen to the message read aloud by Sal.

Completely Local & Private: Just like all of Sal's operations, memory summaries are generated using your local model and stored in your device's secure database. Nothing is sent to any server.

Ephemeral Sight

Sometimes it's easier to show Sal what you're looking at than to describe it. Ephemeral Sight lets Sal take a quick glance at your screen.

How to Use

Instead of pressing  to send your message, press    (Shift + Cmd + Return). Sal will glance at your screen and use what she sees to help answer your message.

You can also trigger Ephemeral Sight with Push to Talk — either force press (Trackpad) or Shift + click-and-hold (Mouse) the Sal orb to dictate your message and let Sal glance at your screen when you release.

Make sure whatever you want Sal to see is visible on screen before you send. There's no visual or audio feedback — she just looks.

Examples

- "What's wrong with this layout?" — while viewing a design
- "Summarise this chart" — while looking at a spreadsheet
- "What does this error mean?" — while staring at a terminal

Setting Up Sight: It is a 2-step process. To enable screen capture, you'll need to grant permission in **System Settings → Privacy & Security → Screen Recording**. macOS requires a one-time restart after enabling.

You're in control. Sal is not recording your screen. There is no video capture, no background monitoring. She takes a single, momentary glance — only when you ask her to — and nothing is stored.

Pinned Documents

Pinning is how you tell Sal: "*Focus on this.*" As long as a document or image is pinned, Sal sees it. It is front and centre for her. Anything you talk about while a file is pinned, Sal will assume you're talking about that item.

How to Pin

1. Drop a file or image into the **input field** — this pins your first document and reveals the pinned documents area.
2. Drop additional files into the **pinned documents area** to pin more.
3. To unpin, remove the item from the pinned area. Sal immediately stops seeing it.

How It Works

1. Sal reads through the document, extracting the important parts.
2. Ask questions about it — Sal will reference what she's read to give you answers.

Why Pinning Is Different

With cloud AI services like ChatGPT, Gemini, or Claude, when you upload a file you don't really know what happens to it. You don't know how long the model sees it. You don't know where it resides on their servers. You have no control over how to permanently delete it. You don't know if it will be used for training data.

We solve all of that with pinning. You know *exactly* what Sal sees, *how long* she sees it, and *when* she forgets it. And because nothing ever leaves your computer, you

know it can never be used for training — by anyone.

Supported file types:

FEATURE	FILE TYPES
Data Intelligence (queryable)	CSV, TSV, XLSX
File Intelligence (readable)	DOC, DOCX, PDF, TXT, MD, HTML, JSON, XML, images, code files
Email Understanding	.eml (drag and drop from Apple Mail)

Email from Apple Mail

You can drag and drop emails directly from Apple Mail into Sal's chat window. She will parse the email, extract the headers and body, and make any attachments available for you to use.

Once pinned, you can ask Sal to:

- Summarize the email thread
- Draft a response
- Extract action items or key information
- Read and analyze any attached files

How to drop: Open Apple Mail, select the email you want, and drag it directly into Sal's input field or pinned documents area. Sal will process it locally — nothing is sent to any server.

Lightning Recall

Lightning Recall is Sal's semantic file search — a way to find files by *meaning*, not just filename. Ask Sal to find something and she'll search your indexed folders using full-text semantic matching.

How to Use

Simply ask in natural language: *"Find the presentation about Q3 revenue"* or *"Where's that contract from Acme Corp?"*

How It Works

- Sal indexes your documents so she can search them by meaning, not just filename.
- Results are scored by relevance and recency.
- Sal can open, read, or act on any result she finds.

The Results Card

When Sal finds matching files, she displays them in a special **results card**. From there you can:

- Click a file to open it directly.
- Click the folder icon to open the enclosing folder in Finder.
- Click the ✕ to dismiss the card when you're done.

What Sal Indexes

Sal does her best to index the things that matter and skip the things that don't — system folders, coding directories, and other technical clutter are automatically ignored.

If you want to explicitly tell Sal to skip something, place it in any folder named **Archive** (or **archive**). Sal will never read or index anything inside an Archive folder. This is a simple way to keep files on your machine without them showing up in Sal's searches.

Web Search

Sal isn't a search engine — she's a researcher. When she needs information beyond what she already knows, Sal will search the web on her own, visit relevant websites, and read what she finds. She uses this to give you answers that are as complete, accurate, and up to date as possible.

You don't need to do anything special. Just ask your question. If Sal needs the web to answer well, she'll go look — and she'll tell you what she found.

Examples

- "What's the latest on the macOS Sequoia release?"
- "How do I fix this Swift compiler error?"
- "What's the weather in San Francisco?"

Privacy-first search. Sal searches through DuckDuckGo, a privacy-respecting search engine. No tracking, no cookies, no user profiling.

URL Context

Paste any URL into your message and Sal will fetch and read the page content. She extracts the main article text, filters out ads and navigation, and distills key facts — all locally.

Examples

- "Summarise this article: <https://example.com/long-article>"
- "What are the key points from this page?"

Goal Seeking BETA

Beta Feature. Goal Seeking is an early-access, autonomous agentic capability. It is powerful but experimental — behavior and output quality may vary. We recommend reviewing all output carefully before acting on it.

Goal Seeking is Sal's most powerful capability. Instead of answering a question in a single response, Sal becomes an autonomous worker — breaking down your goal, managing her own workspace, and executing every step required to complete it. Depending on how you phrase your goal, Sal will direct her effort either outward to the web or inward to your local machine.

Two Modes, One System

Goal Seeking operates in two modes determined by how you phrase your request:

- **Web Research Mode** — Triggered by phrases like *"Deep research [topic]"* or *"Do a deep dive on [topic]."* Sal autonomously browses the web, reads articles, synthesises findings, and compiles a polished report saved to your Mac.
 - *"Deep research Apple's AI strategy."*
 - *"Do a deep dive on the competitive landscape for electric vehicles."*
 - *"Research the history and current state of fusion energy."*
- **Local Action Mode** — Triggered by phrases like *"Goal seeking: [task]"* or *"Seek goal: [task]."* Sal performs autonomous, multi-step actions directly on your machine — organising files, auditing codebases, executing changes — until the objective is complete.

- *"Goal seeking: organise this pinned folder into subdirectories by file type."*
- *"Seek goal: audit this pinned codebase for deprecated APIs and update them."*

How long does it take? That depends entirely on the complexity of your goal. A focused, well-scoped goal might complete in 10–15 minutes. A broad or multi-faceted goal can take an hour or more. We strongly recommend working with Sal in *normal chat first* to refine and pressure-test your goal statement. Once your goal is locked in, copy and paste it directly into the Goal field on the Setup Card.

Step 1 — Triggering a Goal Seeking Job

Sal intercepts your request, extracts a clear **Goal**, and presents you with a Goal Seeking Setup Card before doing anything. Nothing starts automatically.

Step 2 — The Setup Card

Before any work begins, Sal presents a **Goal Seeking Setup Card** with two required fields:

- **Goal Editor:** Sal's extracted goal appears in an editable text field. You can tweak, expand, or completely rewrite it. This is the most important step — a precise, well-scoped goal produces far better results than a vague one.
- **Folder Picker:** Click the folder icon to choose a destination directory on your Mac where Sal will save all output files.

The **Start** button remains disabled (grayed out) until both a Goal and a Folder have been provided. Once both are filled in, the button lights up orange. Click it to begin.

Step 3 — The Active Goal Seeking Experience

Once you click Start, Sal takes over. The UI transforms to give you full visibility into what's happening:

- **The Orange Orb:** Sal's orb indicator turns orange, signaling that an autonomous loop is active.
- **Live Activity Status:** The top of the card displays the current phase — e.g., *"Searching..."*, *"Evaluating Progress..."*, or *"Writing Report..."*
- **Live Thinking Console:** A dark, scrollable console streams Sal's internal reasoning in real-time using a monospaced font. The console auto-scrolls to the latest output.
- **Background Multitasking:** You can minimise the card or click away — Sal keeps working in the background and remains fully available for normal chat. The orange orb lets you know she's still going.
- **Cancel:** A red **Cancel** button is always visible. Clicking it immediately aborts the loop and saves whatever progress has been made to the working log.

Step 4 — Completion & Deliverables

When Sal determines the goal has been met, you'll receive a system notification. At that point:

- The Cancel button becomes a **Dismiss** button.
- An **Open Folder** button appears — clicking it launches Finder directly to your destination directory.

Inside the folder you'll find two files:

- `Goal_Seeking_Log.md` — The raw, chronological record of every action, query, and decision Sal made during the job.
- `Goal_Seeking_Report.md` — The final, polished summary of what was accomplished.

Resuming a Cancelled Job

If you cancel a job or quit Sal mid-way, you don't lose your progress. Sal can pick up exactly where she left off:

1. Locate the `Goal_Seeking_Log.md` file in your output folder.
2. Drag and drop it into Sal's **pinned file area**.
3. The Goal Seeking Setup Card reappears in a **Resume** state, with the original goal and folder pre-populated.
4. Click **Resume** — Sal continues from where she stopped.

If Goal Seeking Triggers by Mistake

Occasionally, Sal may interpret a message as a Goal Seeking request when you intended a normal chat response. If the Setup Card appears unexpectedly, just click **Cancel**. Sal will dismiss the job and immediately re-process your original message as a standard chat query — no need to retype anything.

Fully local & private. All output is saved directly to your Mac. No data is sent to any server. Web searches use DuckDuckGo — no tracking, no user profiling.

Data Intelligence BETA

Beta Feature. Data Intelligence is an early-access capability. We recommend double-checking results against your source file, especially on large or complex spreadsheets.

Sal can read and understand massive spreadsheets using plain conversation — no matter how large the file. Pin a data file to the chat and just ask your question. Sal figures out the answer and gives it to you directly.

This is one of the quiet advantages of running locally. Most AI tools — even the most sophisticated cloud models — read your spreadsheet like a document and make their best guess. Because they process everything as text, they can miss rows, approximate totals, or confidently give you a number that's slightly wrong. Sal works differently: she computes the answer directly from your data, the same way a calculator would. The result is exact, every time, regardless of how large the file is.

Supported File Types

- **CSV** — Comma-separated values
- **TSV** — Tab-separated values
- **XLSX** — Excel spreadsheets (automatically converted to CSV)

How to Use It

1. Drag and drop your `.csv`, `.tsv`, or `.xlsx` file into Sal's chat window, or use the attachment button to pin it.

2. Sal evaluates the file structure to determine whether it qualifies for Data Intelligence (see *File Requirements* below).
3. You'll see a notification telling you what Sal is doing with the file.
4. Ask your question in plain language:
 - *"What is the average sales volume for Q3?"*
 - *"Which product category had the highest revenue last month?"*
 - *"Show me everything where the status is 'pending'."*

File Requirements

Not every file qualifies for Data Intelligence. When you pin a CSV, TSV, or XLSX file, Sal checks whether it meets the following criteria:

1. **Row 1 must be the header row** with distinct, non-empty column names.
2. **All data rows must have the same number of fields** as the header row.
3. **No metadata, preamble, or title rows** above the header.

If Your File Qualifies

Sal loads the file as a queryable dataset and displays a notification:

"Data Intelligence: Analyzing [filename] (N rows, M columns)"

From this point, you can ask questions in plain English and Sal queries the data using SQL behind the scenes — giving you precise, computed answers rather than guesses.

If Your File Does Not Qualify

The file is not rejected — Sal still reads it, just as a document rather than a structured dataset. She'll display a notification explaining why:

"[filename] will be read as a document — [reason]"

Reasons you may see:

- **Header row not detected** — Row 1 is empty, contains data rather than column names, or is missing.
- **Duplicate column names in header row** — Two or more columns share the same name.
- **Inconsistent row structure** — Some rows have more or fewer fields than the header.

If your file is being read as a document and you'd like the full Data Intelligence experience, check the formatting against the criteria above and re-pin the corrected file.

Scope & Limitations

- **Single-file analysis only.** Data Intelligence works on one file at a time. Combining data across multiple spreadsheets is not currently supported.
- **XLSX conversion.** Excel files are automatically converted to CSV before analysis. The same acceptance criteria apply to the converted output.
- **Files with preambles.** Some exports (common in financial reports, lab data, etc.) include metadata or title rows above the actual data. These files will be treated as documents unless the preamble rows are removed before pinning.

Key Benefits

- **Unlimited Scale:** Works instantly on files with hundreds of thousands of entries — large files won't slow Sal down or cause errors.
- **Automatic Data Cleaning:** Sal tidies up date formats and handles blank cells automatically. You don't need to clean or reformat your spreadsheet before asking questions.

Pinning is required. If you type a file path (e.g., `/Users/name/data.csv`) into the chat instead of pinning the file as an attachment, Sal will try to read it as a plain document — which may fail on large files. Always pin the file directly to get the full analysis experience.

Fully local & private. Your data never leaves your machine. Everything Sal learns from your file exists only for the current session and is gone the moment you unpin the file or close Sal.

Audio & Voice Memos

Sal can transcribe, understand, and answer questions about any audio recordings on your Mac. Drag and drop any supported audio file directly onto the Sal chat window, and she will process it locally on your device.

Supported Audio Format

For transcription, Sal accepts `.m4a` files. This format is the standard format used by macOS and iOS for audio recording.

Voice Memos Integration

Voice Memos recorded on your iPhone, iPad, or Apple Watch sync automatically to your Mac via iCloud. This makes them immediately available to use with Sal:

1. Open the native Apple **Voice Memos** app on your Mac.
2. Select the recording you want Sal to analyze.
3. Drag and drop the recording directly from the list in the Voice Memos app into Sal's chat window.

Example Use Cases

- **Meeting & Interview Summaries:** Ask Sal to *"Summarize the key decisions from this recording"* or *"Extract the action items."*
- **Lecture Reviews:** Review a long class or lecture by asking *"What did the speaker say about the project deadline?"*

- **Search and Reference:** Find specific points in a long discussion: *"Search the recording for any mention of the budget numbers."*

Fully local & private. Transcription, speaker separation, and AI summarization happen entirely on your Mac. Your audio files and transcripts are never uploaded or sent to any cloud servers.

Coding & Terminal

Sal can write, execute, and monitor shell commands directly inside your conversation, turning natural language requests into direct terminal actions.

How to Enable the Terminal

To protect your system, terminal execution is **disabled by default**. You must explicitly opt-in to turn it on:

1. Click the Sal icon in your menu bar and open **Settings**.
2. Navigate to **Tools → Terminal**.
3. Toggle the setting to **On**.

Capabilities

Once enabled, you can ask Sal to automate a wide range of developer and command-line tasks, such as:

- **Running scripts:** `python3 main.py`, `npm run dev`
- **Installing dependencies:** `npm install`, `pip install -r requirements.txt`
- **Reading and parsing logs:** Tailing log outputs or reading script errors.
- **Navigating and managing files:** Creating, moving, or editing files and directories.

Command Intelligence Progress Cards

When Sal runs a terminal command, she displays a real-time progress card inline in the conversation. This card shows:

- The active working directory: `[Working directory: /path/to/folder]`
- The exact shell command being run.
- A status indicator showing if the command is running, completed successfully, or failed.

Task Management & Cancellation

- **Cancellation:** You can stop any running command at any time by clicking the **Cancel** button directly on the inline progress card, or by typing `cancel [Task ID]` in the chat window.
- **Background Execution:** Sal can run long-running processes (like local web servers or file watch daemons) in the background. If the Sal app is closed, quit, or crashes, she safely terminates and cleans up all active background processes.
- **Interactive CLIs (PTY):** Interactive commands that require a terminal input or login flow (e.g. `gcloud auth login` or `stripe login`) are wrapped in a pseudo-terminal (PTY) environment automatically.
- **Custom Environments:** Commands are executed in your login shell (`/bin/zsh -lic`), meaning your custom environment variables, paths, and shell tools (like NVM for managing Node versions) are loaded automatically.

Memory Recommendation: Coding and Terminal usage works best on Macs with 32 GB of unified memory or more. On these systems, Sal automatically loads her most powerful 27B parameter reasoning model to analyze and construct commands.

Safety Guardrails: While read-only actions (like checking git status or searching code) run automatically, destructive actions (such as deletes, package installations, or writing to files) always require your explicit approval before running.

Calendar Integration

Sal integrates with your macOS Calendar. She proactively monitors your schedule so she's always aware of what's coming up — without you having to ask.

Capabilities

- **View events:** "What's on my calendar today?"
- **Add events:** "Schedule a meeting with Alex tomorrow at 2pm"
- **Update events:** "Move my 3pm to 4pm"
- **Delete events:** "Cancel the standup on Friday"
- **Open Calendar:** Sal can foreground the Calendar app for you

Proactive Awareness

Sal keeps an eye on your calendar in the background. When your schedule changes, she knows — so she's always up to date without you having to tell her.

Your next upcoming event is displayed in a special **card** in Sal's panel, so you can always see at a glance what's coming up next.

Permissions: On first use, macOS will ask you to grant Sal access to Calendar. This is required for read/write integration.

Reminders Integration

Sal can create, update, delete, and manage your macOS Reminders using natural language — just like she does with calendar events.

The Reminders Card

When you have active reminders, a **Reminders card** appears automatically in Sal's panel. You can:

- Check off reminders directly from the card — no need to open the Reminders app.
- Expand the card to see all your reminders, or collapse it to save space.
- When collapsed, the card shows a badge with the number of urgent or past-due reminders.

Examples

- "Remind me to call the dentist tomorrow morning"
- "What reminders do I have this week?"
- "Mark the grocery list reminder as done"
- "Delete the reminder about the dry cleaning"

Permissions: macOS will ask you to grant Sal access to Reminders on first use.

Email Integration

Sal can compose and send emails through your default macOS mail client. She drafts the content and opens your mail app with everything pre-filled — you review and click Send.

How It Works

Ask Sal to write an email and she'll open the native macOS composer with the recipient, subject, and body already filled in. You can review, edit, and send from your mail app as usual.

Attachments

Sal can attach files to your emails. If you have documents pinned in the chat, she'll include them automatically. You can also specify files by path: *"Email this report to Sarah and attach the budget spreadsheet"*

Examples

- *"Write an email to alex@company.com thanking them for the meeting"*
- *"Send an update to the team about the project timeline"*
- *"Email the contract to legal@company.com with the signed PDF attached"*

You're always in control. Sal opens the composer but never sends emails automatically. You must review and click Send in your mail app — this is a deliberate privacy design.

Image Generation

Sal can generate images right on your Mac. No cloud, no usage limits, no content policies beyond your own judgement. Just describe what you want and she'll create it.

First Time Setup

The image generation model is large (~24 GB) and downloads separately from Sal's core models. On your very first launch, Sal automatically begins downloading it in the background once her core models are ready. This can take 15–45 minutes depending on your connection — but you can start using Sal for everything else immediately. You can also cancel this download if you'd prefer to save bandwidth.

How to Use

Simply describe what you want: *"Generate an image of a cozy cabin in the mountains at sunset"*

Quality Options

Sal automatically adjusts image generation quality based on your Mac's memory:

- **Up to 28 GB RAM:** Fast drafts at 4 steps — ideal for quick iterations and concept exploration.
- **32 GB RAM or more:** Pro-level quality at 20 steps — unlocks high-fidelity, detailed rendering for production-ready images.

You don't need to configure anything — Sal detects your hardware and uses the appropriate quality setting automatically.

What to Expect

Image generation is resource-intensive — your Mac will be working hard while Sal creates the image. She'll display a **progress card** so you can see exactly where she is in the process. The results are worth the wait.

Generated images appear inline in the conversation once complete. Double-click any image to open it in the native Mac Preview app, where you can touch it up, share it, or save it to your hard drive.

Image-to-Image Editing

Sal can modify existing images using text prompts. Pin an image to the chat and describe how you want it changed — she'll generate a new version guided by your instructions while preserving the original composition.

How to Use Image Editing

Drop an image into the pinned documents area, then describe your desired changes: *"Make this photo look like a watercolor painting"* or *"Add a sunset background to this image"*

Control the transformation: Sal automatically adjusts how much of the original to preserve. For subtle edits (color grading, style transfer), she keeps most of the source intact. For major changes (adding/removing objects), she allows more creative freedom.

Updates

From time to time, we release improved versions of the image generation model. When an update is available, Sal will show a notification card at the top of her panel with the new version and size. You choose what to do:

- **Update:** Sal will remove the old version to free up disk space and download the new one. Image generation will be offline for 15–45 minutes while this completes.
- **Later:** Sal will dismiss the notification and won't ask again about that version. Your current image generation continues to work normally.

If you dismiss an update and change your mind later, you can always start it manually from **Settings → Intelligence Engine → Draw Package**.

Note: If you cancel an update while it's in progress, the old version will already have been removed to make room. You'll need to go to **Settings → Intelligence Engine → Draw Package** and restart the download to restore image generation.

About the progress bar: The Draw Package is made up of several large files. The progress bar will reset to 0% as Sal moves between them, so don't worry if it appears to start over — it's making progress. The total download is roughly 24 GB.

Image Understanding

Sal can analyse images — screenshots, photos, diagrams, documents. There are two ways to share an image with her:

- **Pin it:** Drop the image into the input field or the pinned documents area. Sal will see it as long as it's pinned — great for detailed questions or ongoing reference.
- **Ephemeral Sight:** Use   **Return**, force press (Trackpad), or Shift + click-and-hold (Mouse) the Sal orb to let her glance at your screen. Nothing is stored.

Image-to-Image Editing

This feature also works in reverse. You can use pinned images as a starting point for generating new variations. When you have an image pinned and ask Sal to generate something related, she can use it as a reference to guide the output.

Examples

- "What's in this photo?"
- "Read the text in this screenshot"
- "Describe this diagram"

Speech Synthesis

Sal can speak her responses aloud using your Mac's built-in voices — no network required.

What Makes It Different

When Sal speaks, she doesn't just read her response word for word. She actually **summarises it conversationally**, so it sounds natural — like someone telling you the answer rather than reciting five paragraphs of text. This makes listening to Sal feel like a real conversation.

Customisation

You can adjust Sal's voice in Settings:

- **Voice selection:** Choose from any of the voices available on your Mac.
- **Speech rate:** Speed it up or slow it down to your preference.
- **Auto-speak:** Turn this on and Sal will speak every response automatically — no need to ask each time.

Controls

- Press   (Cmd + Period) to interrupt Sal while she's speaking.
- Click the  menu on any message and choose **Speak** to hear a specific message.
- Ask Sal to "read that aloud" or "say that again" anytime.

Things You Can Ask Sal To Do

There's a lot Sal can do for you — all you have to do is ask. Just say what you need in plain language and she'll take care of the rest. Remember, Sal can only help with the things you've given her permission to access.

Files & Folders

- "Open the project proposal on my Desktop"
- "Read the contents of notes.txt"
- "Create a new file called meeting-notes.md"
- "Make a folder called Vacation Photos in my Documents"
- "Move the budget spreadsheet to the Finance folder"
- "Rename that file to final-draft.pdf"

Calendar

- "What's on my calendar tomorrow?"
- "Schedule a meeting with Sarah at 2pm on Friday"
- "Move my dentist appointment to next week"
- "Open the Calendar app"

Reminders

- "Remind me to call the landlord tomorrow morning"
- "What reminders do I have this week?"
- "Mark the grocery reminder as done"

- "Delete the reminder about the dry cleaning"

Web & Research

- "What's the latest news on the Mars mission?"
- "Read this article for me: <https://example.com>"
- "Search for the best Italian restaurants nearby"
- "Research Apple's AI strategy" — starts a Deep Research job
- "Do a deep dive on the competitive landscape for project management software"

Images

- "Generate an image of a sunset over the ocean"
- "What's in this photo?" (with an image pinned)
- "Make this photo look like a watercolor painting" (image-to-image editing)
- "Add a sunset background to this image" (image modification)

Email

- "Write an email to alex@company.com thanking them for the meeting"
- "Send an update to the team about the project timeline"
- "Email the contract to legal@company.com with the signed PDF attached"
- "Summarize this email thread" (with .eml file pinned)
- "Draft a response to this email" (with .eml file pinned)

System

- "Open Safari"
- "Find files related to the Anderson project"
- "Remember this for later"

About approval: When Sal needs to do something that changes your system — like creating a file, moving a folder, or adding a calendar event — she'll ask for your OK before doing it. Read-only actions like reading a file or checking your schedule happen automatically.

Personality Settings

Sal is more than a utility — she has a personality. You can shape how she communicates, sounds, and presents herself. Find these settings under **Settings → Identity**.

Options

- **Name:** Change what Sal calls herself.
- **Persona:** Choose from preset personalities that affect tone, communication style, and how Sal responds.
- **Voice:** Choose which Mac voice Sal speaks with, and adjust the speech rate.
- **Orb Colour:** Customise Sal's visual identity with curated colour themes.

Custom Persona

The most powerful personality option is the **custom persona** field under **Settings → Identity → Personality → Custom**. Here you can type in any personality you want Sal to adopt.

The trick? It works best if you simply **name a person or character**. Just type in a well-known personality — a favourite actor, a fictional character, a historical figure — and Sal will adopt their mannerisms, tone, and way of speaking. It's surprisingly effective and a lot of fun.

Try it: Type the name of your favourite character into the custom persona field and start a conversation. You'll be amazed how well Sal picks up their personality.

Adjusting the Text Size in Sal

If you find the text in Sal too small to read comfortably, or if you prefer a highly compact interface that shows more history at once, you can scale Sal's text in real-time.

1. Click the **Sal icon** in your macOS menu bar and choose **Settings**.
2. Locate the **Identity** section and look for the **Text Size** slider.
3. Drag the slider to the left to shrink text (down to **80%**), or to the right to enlarge text (up to **140%**).
4. As you drag, the sample sentence directly below will scale instantly to show you how your chat will look.
5. **Tip:** To quickly reset your text size back to the standard **100%** factor, simply **click on the percentage number** displayed on the right side of the "Text Size" label.

Push to Talk (PTT)

Think of Push to Talk like an intercom button to your assistant. Sal can only hear you when you're pressing it — otherwise, she can't. No always-on microphone, no ambient listening. This is a deliberate privacy design.

Standard Dictation

Click and hold the Sal orb in the input area (or perform a normal press on a Trackpad). It turns orange, and you can speak your message. Release to send. Sal transcribes your speech and responds in the chat.

Dictation + Ephemeral Sight

To have Sal listen to your message *and* glance at your screen when you release, use the "super intercom" shortcut:

- **Apple Magic Trackpad:** Force press (press harder) on the Sal orb.
- **Apple Magic Mouse (or any mouse):** Hold  (Shift) and click-and-hold the Sal orb.

It's like letting Sal briefly stick her head into your office and see what's on your desk while you're talking to her.

Always intentional. Sal never listens unless you're actively pressing the button. There is no wake word, no background audio processing. You decide when she can hear you.

Panel Sizing & Displays

Sal's physical panel can be quickly resized to fit your workflow by clicking the Sal orb menu icon and selecting from the dimension presets:

- **Panel Width:** Narrow, Medium, or Wide.
- **Panel Height:** Quarter height, Half height, or Full height.

Your specific dimension preferences are permanently remembered across sessions.

Multiple Displays

If you use multiple monitors, you'll notice the Sal orb rests securely in the menu bar across *all* of your screens, but the physical interaction panel is strictly anchored to your **Main Display**.

To move Sal to a different monitor, simply open macOS **System Settings > Displays** and set your preferred screen as the "Main display". Sal will seamlessly migrate over to that screen instantly.

Panel Position

Sal's panel can be anchored to three positions along your screen — giving you full control over where she lives in your workflow.

Changing Position

Click the **Sal orb** in the menu bar, hover over **Panel**, then select **Panel Position**:

- **Left** — Sal docks to the left edge of your screen. Ideal if your primary work area is on the right.
- **Center** — Sal floats in the middle of your screen. Great for focused, distraction-free interactions.
- **Right** — Sal docks to the right edge of your screen. The default position for most users.

Your position preference is saved automatically and restored every time you launch Sal.

Tip: Pair Panel Position with Panel Width and Panel Height to perfectly fit Sal alongside your existing app layout — whether that's a full-screen editor, a split-view browser, or a coding environment.

Menu Bar Dropdown

Click Sal's menu bar icon to access quick controls and preferences. The dropdown gives you direct access to the following:

- **Show Sal:** Quickly open Sal's main conversational panel.
- **Stop Talking** ( ): Immediately interrupt Sal if she is speaking aloud.
- **Check for Updates...:** Manually check to see if a newer version of Sal is available.
- **Submit Bug Report...:** Open the built-in bug reporting tool.
- **User Guide:** Open this very manual.
- **Settings...** ( ): Open the main configuration window (where you can manage Voice Blends, VRAM Token Caps, Permissions, Calendars, and Licensing).
- **Panel Position:** Dock Sal to the Left, Center, or Right of your screen.
- **Panel Width / Height:** Adjust the physical dimensions of Sal's main window.
- **Quit** ( ): Shut down Sal and unload the AI models from memory.

Models

Sal uses several AI models, each specialised for a different job. The core models download during initial setup so Sal is ready quickly. The image generation model downloads on demand the first time you need it.

Core Models (downloaded at setup)

WHAT IT DOES	APPROXIMATE SIZE
Vision-Language Model (Image Understanding & Ephemeral Sight)	~5 GB
Speech Synthesis	~300 MB

On-Demand

WHAT IT DOES	APPROXIMATE SIZE
Image Generation (Draw Package)	~24 GB

All models run entirely on your machine. No cloud inference is ever used for your conversations.

Privacy Guards

Sal was built from day one around a single principle: **your data never leaves your Mac.**

What This Means

- **No cloud inference:** All AI processing happens locally via Apple Silicon.
- **No telemetry:** Sal sends no usage data, analytics, or crash reports to any server.
- **No account required:** No sign-up, no email, no profile. Your license key is all you need.
- **Ephemeral screenshots:** Screen captures are processed in memory and immediately discarded.
- **Local conversation history:** Your chat history is stored on your Mac in a local database. It is never synced or uploaded.
- **Network usage:** The only times Sal accesses the internet are: license activation, model download (once), web search (when you ask), and update checks.

You own everything. Sal has no server-side component. There is no "Sal cloud." If you delete Sal, your data goes with it — because it was never anywhere else.

Updates

Sal checks for updates automatically. When an update is available, you'll see a notification in the menu bar.

How It Works

- Updates are signed and verified before installation.
- You choose when to install — updates are never forced.

Licensing

Your Sal license is a one-time purchase tied to a single Mac.

If you would like to use Sal on an additional Mac, a separate license is required for that machine, as licenses cannot be transferred between computers.

Reporting Bugs

Found something that isn't working perfectly? We want to hear about it, and we've built a beautiful, native tool right into the app to help us fix it quickly.

How to Report

Simply click the Sal menu in your system menu bar and select **Report a Bug**. Fill out the quick form to describe what went wrong.

In the form, you will see checkboxes giving you the option to securely include your **Chat History** (which automatically strips out Personally Identifiable Information / PII), the **App Crash Report**, and the **Server Crash Report**.

We highly recommend voluntarily checking all three boxes when reporting a bug. It securely gives our engineers the exact blueprint we need to instantly isolate the issue, allowing us to patch it quickly so everyone benefits.

Your data is severely protected. Diagnostic logs and chat histories are *never* collected automatically. You must explicitly opt-in to send them using the bug reporter. When you do, they are used strictly by our engineers to fix the bug you reported, and are entirely deleted immediately afterward.

Getting Help

The best way to get help is to **ask Sal** — she knows this manual and can answer questions about her own features.

Other Channels

- **Community:** Join the official [HelloPartner.ai Discord Server](#) to report bugs, request features, and talk directly with the founders.
- **Email:** hello@hellopartner.ai
- **User Manual:** You're reading it — also available as a [downloadable PDF](#).
- **Website:** www.hellopartner.ai

Thank you for choosing Sal. We believe AI belongs to you — on your own computer, in your own home. Without asking permission. Without handing over your data. Without paying rent on your own knowledge. That's why we built her.