

Memory: keep and move your context

DOCUMENT	VERSION	STATUS
SB-GD-131	1.0, October 2026	Public. Checked against the product, October 2026

Who this is for

Anyone who wants SynapBridge assistants to remember durable context, bring facts over from another AI tool, or control who can see what is remembered.

Before you start

- Memory stores short facts, not conversations. A fact has a type, a subject, content and a confidence value.
- Memory lives in Library > Memory. The page heading is "Memory graph". You can also reach it from Settings > Memory, which opens the same page.
- The page has List, Graph and Timeline views. The List view and the fact tools are the core of the feature.
- Importing memory facts from another AI tool needs the Personal plan or above. On the Free plan the request is refused with a plan message.
- Keep your provider's export file ready (see Import below).

How memory works

SynapBridge distills a small number of durable facts, such as preferences, project context and decisions, and stores them in workspace memory. Future chats can show those facts as citations to give the model context. Conversation text is not stored as memory.

Facts get into memory in these ways:

- You add one by hand with New fact.
- Automatic memory from your chats extracts facts from your chat turns. This is a per-user setting (see Admin and privacy controls).

- Memory from meetings and documents extracts facts from meeting transcripts and uploaded documents. This is a workspace setting controlled by the owner.
- You import them from a ChatGPT, Claude or Gemini export.

Each fact shows where it came from: User input, Extraction, Inference, Imported (chat history), Auto-extracted from chat, or Auto-extracted from Compare. The source chip links back to the conversation or import job when there is one.

Create and review facts

- 1 Open Library > Memory.
- 2 Choose New fact.
- 3 Pick a Type: Preference, Context, Entity, Relationship, Note or Decision. A Decision also asks for a few decision details. Prediction facts are meant to be saved from a prediction shown elsewhere in the product with Remember this, so do not create them by hand.
- 4 Enter a Subject (1 to 200 characters) and Content (1 to 4000 characters). Set a Confidence value if you want one.
- 5 Choose Create.

To review what is stored:

- Use the List, Graph or Timeline view.
- Filter by Type and Source, search subject or content, and choose Show archived to include archived facts.
- Select a fact to open its details, connections and audit trail. You can add a link from one fact to another with a short relationship label, such as "supports" or "contradicts".
- Edit a fact in place, or use the row menu for Copy content, Duplicate as new fact, Use in Chat, Archive and Restore.
- Select several facts to archive them together, or to draft a workflow, an agent or a prompt template from them.

The Extraction confidence column shows how sure the extractor was that it read a statement correctly. It is not a judgment of how trustworthy the source is.

Suggested memory

Facts found in a meeting or a chat can appear first as Suggested memory. They are not saved until you choose Add to memory, and only you can see that list. Choose Dismiss to discard one.

Trust level and quarantine

Each fact carries a trust level: Trusted, Verified, Standard trust, Unverified or Quarantined.

- A fact you type yourself starts at Standard trust.

- A fact from extraction, inference or an unknown source starts at Unverified.
- A derived fact never gets a higher trust level than the facts it came from.
- A Quarantined fact is hidden from chat, search and retrieval for everyone, including its author, until an owner or admin reviews it.

Owners and admins can review quarantined facts at Settings > Memory quarantine, which the Settings menu shows to admins on the Team plan and above. Restore sets a fact back to Unverified, not to a higher level. The same page lists write-pattern anomalies, which are detection only by default: nothing is quarantined until an admin acts on it.

Who can see a fact

- Each fact is either private to its author or visible to the workspace.
- A fact you add by hand with New fact is visible to the workspace. Facts from extraction and imports start as private to the person whose activity produced them.
- The New fact form does not offer a visibility choice. The author of a fact can change it afterwards from the row menu with Share with workspace or Make private. Cards show Only you or Shared. A fact that is no longer available to you may show as unavailable if it was archived, removed or is private to someone else.
- You can edit or archive facts you own. Owners and admins can also edit or archive other people's workspace-shared facts. They cannot see or change another member's private facts.
- Archiving moves a fact out of active recall for the workspace. Archived facts stay visible when Show archived is on, and you can Restore them.
- Reset all memory permanently deletes only your own facts in this workspace. It has no undo. You must type RESET to confirm. Nobody can use it to delete another member's facts.

Import from another AI tool

There are two importers. They do different things.

Import memory facts (preview, then approve)

Use this to turn an export from another assistant into reviewable facts. Open Library > Memory and choose Import, or open Settings > Import AI history.

- 1 Choose a source: ChatGPT, Claude or Gemini. Perplexity is listed as best effort, with no confirmed export format, so an upload may fail. Slack and Microsoft Copilot are shown as not yet supported.
- 2 Get your export from the provider:
 - ChatGPT: Settings > Data Controls (downloads a .zip).
 - Claude: Settings > Data Privacy Controls (downloads a .zip).

- Gemini: Google Takeout > My Activity > Gemini Apps, with JSON as the format.
- 1 Upload the archive as a .zip file and choose Upload and preview.
 - 2 Wait while the job moves through reading conversations and distilling facts.
 - 3 When the preview is ready, review the sample of candidate facts. Choose Approve and import, or Cancel.
 - 4 When the import completes, the page reports how many facts were imported into your memory graph.

Nothing is added to your memory until you approve the preview. The conversations themselves are not saved by this importer, only the distilled facts.

To undo an import, open the completed import under Settings > Import AI history and choose Undo import, then Yes, archive. This archives all the facts that import added. You can also choose View imported facts first. Undo works only on a completed import, and archived facts can be restored from Library > Memory.

Imported facts are private to you, and facts with identical content are skipped as duplicates.

Import from another AI tool in Chat (conversation titles)

In Chat, the "Your other AI tools" section has an Import from another AI tool button. It reads a single .json file that you have unzipped from the provider's export.

- Supported tools in the picker: ChatGPT, Claude, Gemini, Mistral, xAI, Perplexity, DeepSeek and Qwen. The export file layouts for Mistral, xAI, Perplexity, DeepSeek and Qwen are not verified, so an upload from those tools may not be recognized.
- The file must be unzipped JSON, no larger than 25 MB. The product reads up to 2000 conversations from it and tells you if it stopped early.
- It records conversation titles and, as a best effort, project groupings. It does not create memory facts.
- Nothing is fetched from any provider automatically. You upload the file yourself.

Admin and privacy controls

- Automatic memory from your chats: a per-user switch on the Data & Privacy settings page (/settings/data/privacy). The Settings menu lists Data & privacy only for the Enterprise plan, so open the page by its address on other plans. The switch is on by default. Turning it off stops extraction for your own turns only.
- Memory from meetings and documents: a workspace switch, owner only, on the same page. Turning it off stops new facts from landing this way. Existing facts stay.

- Memory provider rights: owners choose which connected AI providers and custom MCP servers may use workspace memory, at Settings > Memory provider rights (owner only, shown on the Team plan and above). A provider with no setting is blocked.
- Memory quarantine: owners and admins review quarantined facts and write-pattern anomalies (see above).
- Audit: creating, editing and archiving facts is audited, and each fact has an Audit trail tab.

Limits and notes

- Fact limits: subject up to 200 characters, content up to 4000 characters, extra metadata up to 4096 bytes.
- Storage limit by plan for Memory Graph: Free 100 MB, Personal 2 GB, Pro 10 GB, Team 50 GB, Business 500 GB, Enterprise custom. When you reach it, new facts are refused until you free space or upgrade.
- Memory import archive: up to 500 MB. Each import job considers at most 500 candidate facts. The default is 5 import jobs per user in a rolling 24 hours.
- Chat sidebar importer: unzipped .json, up to 25 MB.
- Memory is separated per tenant using row-level security on the memory tables, with the tenant also checked in each query. This document makes no claim about separate encryption keys per tenant.
- Memory facts are one tenant's data. They are not shared with other tenants.
- A plan listing on the pricing page names more import sources than the importer supports. Use the sources named in this guide.
- Documents and files belong in Library > Knowledge, not Memory.