

Responsible AI use in your workspace

DOCUMENT	VERSION	STATUS
SB-GD-143	1.0, October 2026	Public. Checked against the product, October 2026

Who this is for

Workspace owners, admins and team leads who set the ground rules for how people use AI in SynapBridge.

Before you start

- Decide who owns AI rules in your organization. Most controls below are admin settings.
- Decide which kinds of data may go to which models. Write it down in plain terms your team can follow.
- This is practical guidance, not legal advice. Talk to your counsel for legal or regulatory questions.
- Do not enter protected health information (PHI). SynapBridge's HIPAA business associate agreement program is planned. Until a signed agreement exists with you, treat the workspace as not approved for PHI.

Start with policies

SynapBridge applies workspace policy to requests made through SynapBridge. Usage outside SynapBridge is not covered.

- Use the content policy settings to set rules for what is allowed.
- Use model entitlements and connector settings to limit which models people can reach.
- Versioned policies let you review changes and test a rule with the policy simulator before you enforce it.
- Review retention settings so stored content is kept only as long as you need it.

Use sensitive-data detection, and know its limits

SynapBridge scans content for known patterns, such as secrets and credentials, national ID numbers, payment card numbers, and email or phone details. For each detector, an admin policy sets an action: block, flag, redact or allow.

- The default policy blocks several high-risk patterns (for example US Social Security numbers and credit card numbers) and flags lower-risk ones.
- Redaction exists but is off by default. Rewriting a prompt can change the answer, so decide deliberately before you turn it on.
- Detection looks for known patterns. It will miss data that does not match a pattern, and it can flag text that is not sensitive. It is a safety net. It cannot promise that no sensitive data leaves your workspace.

Good habits for your team:

- 1 Do not paste passwords, keys or customer records into prompts.
- 2 Remove names and identifiers when they are not needed.
- 3 Treat a block or flag as a prompt to stop and check, not an obstacle to route around.

Approvals (Beta)

SynapBridge includes approval workflows, including two-person approvals for sensitive changes. This area is Beta. Confirm in your own workspace that an approval step is active before you rely on it as a control.

Keep an audit trail

The Audit page shows tenant events such as security, redaction, policy and cost events. It is available to admins and above. Review it on a regular schedule, and look at what was blocked, flagged or changed.

If you use an external log tool, check your plan for audit streaming and export options.

Review memory

SynapBridge Memory keeps facts that help later conversations.

- Review what is stored. Remove facts that are wrong, outdated or should not be kept.
- Admins can quarantine a fact. A quarantined fact is hidden from chat, search and retrieval until an admin restores it.
- Candidate memories proposed from meeting notes wait for you to confirm or dismiss them.

Check AI output

- AI can be wrong. Check facts, numbers and citations before you act on them.
- Keep a person responsible for decisions that affect other people.
- Tell colleagues and customers when AI helped produce work, where your own rules or agreements require it.

Limits and notes

Note SynapBridge is architected to align with the SOC 2 Trust Services Criteria.

- Controls apply to requests made through SynapBridge. They do not cover tools used outside it.
- Availability of some controls depends on plan. Check your plan page.
- Nothing here is legal advice or a statement of regulatory status.