

Understanding AI cost in SynapBridge

DOCUMENT	VERSION	STATUS
SB-GD-113	1.0, October 2026	Public. Checked against the product, October 2026

Who this is for

Workspace admins, finance leads and FinOps owners who need to read, forecast, attribute and export AI spend.

Before you start

- Most of the cost views below are in the Insights area and are limited to Owners and Admins. Members can see their own calls in the Request log.
- Availability of a view or export depends on your plan. Where a feature is not included, the product shows an upgrade message instead of data.
- Cost views read the cost records SynapBridge keeps for AI calls that go through SynapBridge. They do not read your provider's billing system, except in the optional "Reported directly by provider" section described below.
- Some Insights pages, such as Cost attribution, Budgets & forecasting and Anomalies & alerts, show a badge near the top. "Live data" means the page is showing your records. "Demo data" means the page is showing a labeled sample because a data source is not available. Do not use Demo data for decisions.

What you pay for

- SynapBridge is a subscription. Free, Personal and Pro are flat monthly plans. Team, Business and Enterprise are priced per seat. Current figures are on the pricing page.
- The plans use your own AI provider API keys. Your provider bills you directly for the model usage on those keys. That charge is separate from your SynapBridge subscription.
- Each plan lists a Memory Graph storage allowance on the pricing page.

How cost per response is shown

- A response in chat and Compare shows a cost next to it. When SynapBridge has an exact rate for the model, the cost is calculated from token counts at that rate.
- When SynapBridge has no exact rate for a model, it applies a provider-level default rate. The figure then appears with a leading ~ and the suffix "est.". Hover over it to see the date the rate table was last updated.
- Treat every per-response figure as an estimate. Your provider's invoice is the source of truth.

Read your spend

- 1 Open Insights and find the "Spend & value" group. Usage & spend and Request log are open to every role; most other pages in the group are for Owners and Admins.
- 2 Open Spend & cost analysis for provider and model breakdowns and unit-cost drift.
- 3 Open Usage & spend for token counts and cost by Provider, Model, User or Day. Pick a range: 7d, 30d, 90d, MTD or YTD. Select Export CSV to download the table. If you do not see Usage & spend in the menu, open /insights/usage.
- 4 Open Request log to see individual calls with tokens, cost and latency. Members see only their own calls.
- 5 Open Anomalies & alerts to see spend anomalies and budget threshold events.

On Usage & spend:

- Month to date, Previous month and "On-pace projection" are shown. The projection is month-to-date spend divided by days elapsed, multiplied by days in the month. It is a straight line with no seasonality.
- A query returns a limited number of groups. When the limit is reached, the page says lower-spend groups are not in the table.
- Calls made to a provider without going through SynapBridge do not appear in the routed figures.
- A section labeled "Reported directly by provider" shows usage pulled from your own provider account, once usage tracking is set up for that provider. It uses a separate read-only key and only some providers are supported. It is visible to Owners and Admins. The figures are the provider's own and may lag, and they overlap with the routed figures. Compare them and do not add them together.
- A separate section shows Bring-Your-Own-Model spend, and another shows usage reported by CLI and agent tools. Each is its own series. Do not add them to the others.

Forecast

- 1 Open Insights, then Budgets & forecasting.

- 2 Find the "Forecast to budget" tile. Choose a horizon: 7, 14, 30, 60 or 90 days.
- 3 Read the p50 forecast, the p95 forecast and the Trend (Rising, Steady or Falling).

How to read it:

- The forecast is a linear regression over the last 90 days of cost records. It is fit only on days that had events.
- p50 is the point estimate. p95 adds a buffer for how variable your spend has been.
- A forecast is a planning figure. It is not a bill and not a commitment. It will be wrong when usage patterns change.
- The same page shows a forecast by provider, and a spend cap forecast tile that says plainly when no cap is configured.

Attribution

- 1 Set attribution rules so spend can be assigned to teams, projects and cost centres. Open the rules page at /settings/attribution (Owners and Admins). The menu entry named Cost attribution under Admin center opens the Insights report instead, so use the address. The rules page has Rules, Projects and Summary tabs, plus Recurring and Cost centers.
- 2 Open Insights, then Cost attribution.
- 3 Choose the window, then choose By: User, Team or Project.
- 4 Choose the Framing: Showback or Chargeback. Showback is informational. Chargeback labels the totals as billable to the owning team or project.
- 5 Select a top driver to open its detail.

Provider and model attribution is read from the call itself. User, team and project attribution depends on your rules. If no rules exist, team and project views can be thin or empty.

Chargeback export

- 1 On Cost attribution (or Budgets & forecasting, or Anomalies & alerts), find the Chargeback export panel. Cost attribution also has an "Export invoice" panel that makes an informational PDF invoice for one team, project or user. It is not a billing charge.
- 2 Set the Start date and End date.
- 3 Choose the Format: CSV or JSON.
- 4 Under Group by, pick at least one of User, Team, Project, Provider, Model, Category.
- 5 Select Generate export.
- 6 Select Download. The download link is valid for 15 minutes. Generate a new one if it expires.

Every row includes GL code, GL account and GL description columns. Owners can maintain the GL mapping on the Cost attribution page.

Saved views and weekly schedules

- 1 On the same pages, find "Saved views and scheduled reports".
- 2 Enter a View name, set the dates and group-by choices, then select "Save this view".
- 3 In "Your saved views", select Apply to reload a view, or Schedule to email it.
- 4 Choose a Format (CSV or PDF), the Day and the Hour (UTC), then select "Create schedule".

A schedule is weekly. The email goes to your own address by default. The summary line shows the day, the hour in UTC, the format and the address.

Limits and notes

Note Costs in SynapBridge are records and estimates. Reconcile them against your AI provider's invoice before you rely on them for billing or accounting.

- SynapBridge does not set provider rates. The rate table can lag a provider's price change. The tooltip on an estimated figure shows when the rates were last updated.
- Spend from calls that do not pass through SynapBridge is not in the routed figures. Use the "Reported directly by provider" section where your provider supports it.
- Forecasts use at most 90 days of history and a straight-line model.
- Weekly scheduled emails depend on the scheduled reports service being enabled for your deployment. If an email does not arrive, contact SynapBridge support.
- Cost caps and spend thresholds are separate from these views. See the "Cost caps and alerts" guide. A cap is a check and not a guarantee.
- Do not rely on any cost view for a statement of what you owe. This guide makes no claim about savings.