Fabric Plan

E1 Plan, Calculate, and Optimize

This FAQ addresses common questions and clarifications that arise while working through the E1 — Plan, Calculate, Optimize exercise. It covers environment setup, interface navigation, top-down and bottom-up planning mechanics, the Optimizer, and the approval workflow.

Part 1 Set Up the Training Environment

1. What permissions do I need to create a Plan app?

You need Contributor or higher access to the Fabric workspace and at least Build access to the semantic model you want to connect.

2. Can I connect one Plan app to multiple semantic models or change the semantic model after the plan app is created?

No — each Plan app connects to one semantic model, and it cannot be changed once it is connected to a Plan app. If you need to plan against different data sources, you will have to create separate Plan apps.

3. What happens if my semantic model is not visible in the connection dialog?

Check that you have at least Viewer access to the workspace where the semantic model is published — the semantic model doesn't need to live in the same Fabric workspace as your report or item, since it can be accessed across workspaces. Still, you do need at least viewer-level permissions on its host workspace to see it in the dialog. If it still doesn't appear, try refreshing the connection or signing out and back in.

4. As a contributor, can I create my own semantic model connection?

No, contributors cannot create a semantic model connection directly. If you only have contributor access, the workspace admin should create the connection, and the contributor can use it.

5. What is the Fabric SQL database used for?

The Fabric SQL database is where Fabric Planning stores all plan data — inputs, forecasts, scenarios, and writeback entries. It is separate from the semantic model so that plan data can be written back without affecting the underlying actuals or reporting layer.

6. What's the difference between the semantic model connection and the SQL database connection?

The semantic model connection pulls your actuals/reporting data in for planning, whereas the SQL database connection is a separate connection that enables collaboration — it's what allows the plan to be saved and shared and is set up after the planning sheet is built.

7. Can an existing connection be reused instead of creating a new one?

Yes. In the connection dropdown, an existing connection can be selected instead of choosing Create a new connection, provided it meets the criteria outlined in the prerequisite document.

8. Does the Plan app need to be created before or after the SQL database?

The SQL database is created first in this flow, but the two are independent items — creating one doesn't depend on the other. However, the SQL database should exist before writeback is attempted, since planning data can only be written back to a Fabric SQL database, not any other data source.

Part 2 Explore the Interface

9. What is the difference between semantic model data and From Sheets data in the Data pane?

Semantic model data is read-only actuals and reference data pulled directly from your Power BI dataset — historical revenue, dates, and product hierarchies. From Sheets data is plan data created within Fabric Planning (forecasts, targets, sales plans) and made available for use in other sheets.

10. Can I add dimensions from different tables to the same planning sheet?

Yes — as long as the tables are related in the semantic model, you can combine dimensions from different tables in the rows, columns, and values of a single planning sheet. If the tables are not related, the values will not cross-filter correctly.

11. If I change the field assignments on a planning sheet — for example moving a dimension out of Rows or removing a measure from Values — does it affect

the underlying data?

No — field assignments are a display configuration only. Adding, removing, or rearranging what appears in Rows, Columns, and Values on a planning sheet does not alter the data in the semantic model or the Fabric SQL database. It only changes what is shown on that sheet.

12. What is the difference between a visual-level filter and a global filter?

A visual-level filter applies only to the current planning sheet and does not affect other sheets in the Plan app. A global filter applies across all sheets. In E1, the year filter applied in the sidebar is a visual-level filter — it scopes the Gross Profit sheet to 2025 data without affecting the Plan Intro sheet.

13. How do I navigate a sheet with a large number of rows?

Use the footer's Settings control to configure how many rows are displayed per page, then use the page navigation controls — first, previous, next, and last — to move between pages. For columns that extend beyond the visible area, use the horizontal scrollbar at the bottom of the canvas or reduce the zoom level using the zoom slider in the footer. Collapsing the Explorer pane and sidebar panels also gives the canvas more space.

14. Can I export to PDF or CSV as well as Excel?

Yes — the Export dialog offers PDF, Excel, and CSV. PDF produces a fully formatted, paginated export — preserving layout, cell values, number formatting, notes, filters, and calculated rows and columns. Advanced options let you configure page size, orientation, content order, scaling, margins, and whether to include comments, filter information, or a logo. Excel export is for further analysis and data manipulation, with three export modes: Fully Expanded (all hierarchy levels expanded flat), With Expand/Collapse (hierarchy grouping controls preserved in Excel), and Current State (exports exactly what's currently visible on screen, including applied filters and expanded levels). CSV is a flat data export useful for downstream processing, but it does not preserve hierarchy, formatting, or layout.

15. What actions can be performed and saved while in Reading Mode?

Reading Mode isn’t purely static. While it locks down the core design of your report, it keeps the planning sheet highly interactive for day-to-day work. However, there is a clear distinction between temporary visual changes and saved planning actions.

1. Temporary Actions (Not Saved)

You can customize your view on the fly by changing sheet layouts, updating number formatting, applying Top N rules, or filtering data. However, these visual adjustments are strictly temporary and will not be saved when you exit or refresh the report.

2. Saved Actions

The following core planning and data entry actions are fully saved and committed while working in Reading Mode:

- ✓ Data Inputs: Type numbers directly into cells, as well as bulk edit, allocate, and distribute values across rows and columns.
- ✓ Forecasts: Change planning assumptions and distribute the updated impacts across your forecast columns.
- ✓ Comments: Add, update, and manage commentary.
- ✓ Writeback: Commit your data changes back to the backend database—provided an active writeback destination has already been configured by the report author.

Part 3 Top-Down and Bottom-Up Planning

16. Why do we use 2025 actuals as the basis for the 2026 Target column rather than starting from zero?

Starting from prior year actuals gives the target a realistic baseline — the distribution reflects where revenue actually came from, not an arbitrary equal split. It also makes the growth assumptions visible: if Beer was 85% of revenue in 2025, it should receive roughly 85% of the target unless there is a deliberate strategic reason to shift the mix.

17. What does “Copy from” mean when inserting a column?

"Copy from" pre-populates the new column with values from an existing measure at the time of insertion. It is a one-time copy — the new column is independent afterwards and will not update automatically if the source measure changes. In E1, we copy from 2025 Gross Revenue to create a starting baseline for both the 2026 Target and 2026 Sales Plan columns.

18. Can I edit or delete a Data Input measure/column?

Yes. Click the gripper icon on the column and select Edit/Delete Measure. Alternatively, click on

Manage Measures in the Planning ribbon and edit or delete the measure/column from the side panel.

19. What happens to the target values if I change the planning hierarchy later — for example, adding a new brand?

New rows added to the hierarchy after the target has been set will show blank values in the target column. You would need to manually enter or distribute a value for the new row, which will aggregate up and change the parent and grand totals.

20. What is the difference between locking a cell and locking all children?

Locking a cell prevents that specific cell's value from changing when a parent or sibling value is updated. Locking all children locks every descendant of the selected cell — useful when you want to protect an entire subtree. In E1, we lock the Q1 and Q2 grand totals so that when we re-enter $18.75B, the remaining $0.25B distributes only across Q3 and Q4.

21. Does keeping the cells locked hide them when the measure is pulled via From Sheets in another planning sheet?

No — locked cells in a planning sheet are local to that sheet's editing session. If the target values remain locked and another sheet tries to read from them via From Sheets, the values are still accessible — locking does not hide data.

22. How do I unlock a locked cell, and what happens when I do?

Click the locked cell and select Unlock from the cell's context menu. Unlocking doesn't change the cell's current value; it simply makes the cell eligible for redistribution again the next time a parent or sibling value is updated.

23. What does Distribute to column with trend do?

It takes the value in the selected cell and spreads it across the column hierarchy following a growth curve rather than evenly. In E1, after appending Summit Lager's Grand Total by 8%, we apply a 4% trend to distribute that annual total across quarters — Q1 receives the smallest share and Q4 the largest, reflecting a pattern of growing demand through the year.

24. In Bulk Edit, what's the difference between Apply to Row Level and Apply to Column Level?

Row Level determines which row granularity the edit is calculated and applied at (e.g., Brand), and Column Level determines which column granularity it's applied at (e.g., Month). Together with the filters, they scope exactly which cells the bulk change touches.

25. Does Bulk Edit overwrite existing values?

It depends on the Type selected in the Bulk Edit dialog: Set Value replaces the existing value with the number you enter; Append By adds or subtracts from the existing value; Reset discards any plan inputs and restores the original seeded values.

26. Can Bulk Edit apply different percentages to different rows in the same operation?

No — a single Bulk Edit operation applies one value or percentage to all rows matching the filter criteria. If you need different percentages for different rows, run separate Bulk Edit operations with different filters.

Part 4 Optimizer

27. Why create separate Sales Plan and COGS input columns instead of using the original measures directly?

The Optimizer requires editable data input columns to overwrite the optimal values. The original 'From Sheets' Sales Plan and the 2025 COGS semantic model measure remain unchanged and are kept as a baseline for comparison after the Optimizer runs.

28. What do Column aggregation: Formula and Row aggregation: Formula mean when creating Gross Profit?

They tell the sheet to calculate Gross Profit using the formula at every column and row level (e.g., each month and each brand), rather than summing pre-calculated values. Optimizer requires both row and column aggregation to be set to formulas.

29. What is the purpose of adding constraints to the optimizer?

Constraints are used to establish a range within which the variable will have to fall - for example, limiting COGS reduction to a maximum value.

30. Can measures other than Sales Plan and COGS be selected as variables?

Only editable data input columns on the sheet can be selected as variables. Measures pulled from From Sheets or the semantic model — like the original 2026 Sales Plan or Total COGS — are not editable and cannot be chosen as variables.

31. What is the difference between the Objective Target and a Constraint?

The Objective Target defines the single value the Optimizer is solving toward — Gross Profit = $12.5B in this exercise. Constraints are additional boundaries placed on the variables that the solver must stay within while finding a solution — for example, capping how much COGS can be reduced.

32. What happens if the $12.5B target genuinely cannot be reached?

The Output screen shows the Achieved value alongside the Target Value. If the solver cannot fully close the gap, Achieved will differ from Target and a red indicator appears — signaling the result is the closest feasible solution, not an exact match.

33. Does the Optimizer respect locked cells?

Yes — if a cell is locked, the Optimizer's redistribution skips it entirely. Its value remains unchanged even as the rest of the hierarchy is adjusted to hit the target.

34. After clicking Apply, can the optimized values be reverted?

Apply writes the Optimizer's solution directly into the Sales Plan and COGS columns, overwriting the prior values. There is no built-in undo action. Reverting requires manually re-entering the original figures.

35. Can the Optimizer be re-run with a different target or constraints without rebuilding the Gross Profit formula?

Yes — Gross Profit is a formula column that recalculates automatically from the Sales Plan and COGS inputs. The Optimizer can be launched again from the Grand Total cell with new objectives or constraints without recreating the formula.

36. When would you use Optimizer over a normal top-down planning approach?

Top-down planning works by manually setting a target at the top and allocating it down to lower

levels — it's straightforward, but the allocation is based on whatever logic the planner chooses, not on what's actually achievable given real constraints. Optimizer is used when the target itself needs to be solved for, given a set of constraints — it calculates the combination of values thatbest satisfies the goal, rather than having the planner distribute a number manually. So top-down planning suits simple, judgment-based allocation, while Optimizer suits scenarios where multiple constraints need to be balanced to arrive at the optimal plan.

Part 5 Approvals

37. Who receives a Teams notification when a row is submitted?

If Teams Notifications is enabled during setup, the person whose name or email was entered for the Finance Director level receives a Teams notification once a row is routed to them for approval.

38. What does Reset on Rejection do?

When enabled, a rejected submission resets its approval status, requiring the submitter to resubmit after making changes rather than leaving it in a rejected state indefinitely.

39. Does the Finance Director/Approver get notified automatically?

Yes, if Teams Notifications are enabled during setup. A Teams notification is sent to the approver once a row is routed to them.

40. Do Teams Notifications require the approver to have Microsoft Teams access?

Yes — notifications route through Teams specifically. There is no automatic email fallback built into this setting.

41. Can the Approval workflow be deleted after it has been configured?

Yes — navigate back to Approval and eject the flow. Once ejected, the Regional Manager and Finance Director columns can be deleted from the sheet.

42. Are two levels the maximum, or can more approval levels be added?

Two levels reflect the scenario in this exercise, not a fixed cap. Additional levels can be configured, each adding its own approval column to the sheet.

43. Does the Approval workflow apply to the whole sheet or just the 2026 Sales Plan column?

It is configured at the sheet level, but the approval columns track status per row rather than gating every column. In this exercise, it governs changes reflected through the Regional Manager and Finance Director columns tied to the 2026 Sales Plan edits.

44. Once the Finance Director column shows Pending, can the Regional Manager still edit the underlying value?

Yes — Pending reflects the approval routing status, not a lock on the cell. The underlying 2026 Sales Plan value remains editable. The workflow tracks and displays approval progress but does not restrict data entry on the sheet.