

E2**Forecasting and Writeback**

This FAQ addresses common questions and clarifications that arise while working through the E2 — Forecasting and Writeback exercise. It covers rolling forecast concepts, statistical forecasting configuration, closing and extending the forecast horizon, reforecasting open periods, and writing the finalized forecast to the Fabric SQL database.

Part 1**Create the Rolling Forecast****1. What is the difference between a closed period and an open period in a forecast measure?**

A closed period reads from an existing measure — typically actuals — so the forecast reflects what actually happened. An open period accepts data input, either manual entry or values copied from another measure during period setup. As time moves forward, you progressively close periods to replace forecast values with actuals, keeping the full-year view accurate.

2. Why does the forecast measure start with blank values in E2 rather than being seeded from 2025 actuals?

The blank starting point is intentional — it ensures that the Statistical Forecast model populates the 2026 values entirely from the historical pattern it detects in 2024 and 2025 actuals, without any manual assumptions influencing the starting point. If the measure were pre-seeded, the statistical model would be adjusting existing values rather than generating them from scratch.

3. Can I have more than one forecast measure on the same planning sheet?

Yes — you can create multiple forecast measures and display them as separate columns on the same sheet. This lets you compare the original forecast against a reforecast over time. Each forecast measure has its own open and closed period configuration.

4. Can I delete a forecast measure and start over?

Yes — go to Planning → Manage Measures to delete the forecast measure. This removes the measure and all associated open period configurations from the planning sheet.

Part 2**Generate Statistical Forecast**

5. Why do I need at least two years of historical data for the statistical forecast?

The algorithm needs enough data to detect seasonality patterns reliably. With less than two years, the model cannot distinguish a true seasonal pattern — for example, a Q4 uplift every year — from a one-time event. Two years gives the model one full seasonal cycle to learn from and one to validate against.

6. What does the Confidence % setting control?

The confidence interval defines the range within which the algorithm expects the true value to fall, given the variability in historical data. A 90% confidence interval means the algorithm is 90% confident that the actual value will fall within the shaded range shown in the forecast preview. A higher confidence produces a wider range; a lower one produces a narrower range. The point estimate — the central line — does not change; only the uncertainty band changes.

7. What does the Growth Factor of 4% do — is it applied on top of the model's prediction?

Yes — the Growth Factor applies an additional uplift on top of the model's statistically derived values. Setting it to 4% instructs Predict to grow its output by 4% before saving, reflecting a business assumption about market growth that the historical data alone wouldn't capture.

8. What is the difference between Bottom-Up and Top-Down evaluation?

Bottom-Up calculates the forecast at the most granular level first — brand level in E2 — and aggregates up to category and total. Top-Down calculates at the top level and distributes down. Bottom-Up is generally more accurate for granular sales data because each brand's seasonality pattern is captured independently. Top-Down is faster but smooths over brand-level differences.

9. If Bottom-Up runs at the Brand level, why do I click the Grand Total cell to start the run?

The Grand Total cell is just the trigger point — it tells Predict which measure to populate. The Bottom-Up evaluation method internally iterates through each brand and aggregates up; clicking the Grand Total cell doesn't restrict the model to only the Grand Total row.

10. What is the difference between the available forecasting algorithms?

Three algorithms are available under Choose Algorithm. Trend Decomposition with MSTL splits the series into separate layers — one per seasonal cycle plus a trend — forecasts each layer independently, then combines them. It is the best choice when data has patterns repeating at multiple levels, such as both annual and quarterly cycles in beverage sales. Exponential Smoothing predicts by weighting recent observations more heavily than older ones — a good general-purpose choice for most business data including sales and demand. ARIMA models the statistical structure of the series based on how each value relates to past values and past forecast errors — more technical and better suited to data with strong autocorrelation patterns.

11. What does the Set Seasonality setting under Customize Algorithm do?

Set Seasonality tells the model which cycle lengths to detect and account for when generating the forecast. Ticking Year instructs the model to look for patterns that repeat annually — for example, a consistent Q4 uplift each year. Ticking Quarter adds a second layer, capturing patterns that repeat within the year at the quarterly level.

12. What does the forecast preview show before saving?

Historical data is shown in grey, while predicted values are shown in green, with the confidence range displayed as green shading in the background. This allows you to visually assess whether the model's output looks reasonable before committing the values to the forecast measure.

13. Can I re-run the statistical forecast with different settings after saving?

Yes — you can re-run Predict on the same measure with different settings. The new values will overwrite the previously saved forecast values for the selected date range.

14. Can I apply the statistical forecast to a subset of categories or brands rather than the full hierarchy?

Yes — you can filter the rows before running Predict, or configure the row selection in the Predict dialog to target a specific subset. The Bottom-Up evaluation will then run only for the rows in scope.

Part 3

Close the Period, Extend, and Reforecast

15. What exactly happens when a period is formally closed?

The closed month is locked — it populates automatically with the linked measure and turns grey to indicate it cannot be edited. Any forecast assumption previously entered for that month is replaced by the actual value.

16. Can I edit the closed period values?

No — closed period values cannot be edited. They will be greyed out, indicating that those values are non-editable.

17. What does “Overwrite forecasts” do versus “Retain forecasts” when closing a period?

Overwrite forecasts replaces the closed period forecast values with actual performance data — this is the default and the most common choice for a rolling forecast, since you want actuals to be the single source of truth for closed periods. Retain forecasts keeps the original forecast values and only fills in actuals where forecast values are blank. Use Retain when you want to preserve the forecast for variance analysis alongside actuals.

18. If I close January 2026 and the actuals are slightly different from my forecast, does the grand total change?

Yes — closing a period replaces the forecast value with the actual value, so the full-year total will reflect the actual January figure. This is expected and correct behavior. The variance between forecast and actuals is visible when you show both measures side by side.

19. Can I close multiple periods at once, or only one at a time?

You can close a custom range — not just the previous month. In the Close Period dialog, select Custom under Close Period Till and specify the end date of the range you want to close. This is useful if you are catching up after a gap or closing an entire quarter at once.

20. What does Extend Forecast Range do when closing a period?

Extend Forecast Range automatically adds one period to the end of the forecast horizon when you close a period. This is what makes it a true rolling forecast — closing January 2026 and extending by one month adds January 2027, so the forecast always covers the same forward-looking window.

21. What happens if I close a period without ticking “Extend Forecast Range”?

The period closes, and actuals populate correctly, but the forecast horizon shrinks by one month. You would need to manually extend it later using Reforecast.

22. After closing January, can I reopen it if there is a correction to the actuals?

Reopening a closed period is not a standard operation — it would require adjusting the close period configuration. In practice, actuals corrections are typically handled in the source system, and the close period is re-run.

23. What is the difference between reforecasting and editing the forecast measure directly?

Reforecast opens a new period setup dialog that lets you reconfigure the open period behavior — for example, changing the copy source or growth assumptions — and then redistributes values across the remaining open periods accordingly. Editing cells directly is faster for a single adjustment, but does not recalculate the full distribution. Use Reforecast when assumptions change broadly; use direct editing for targeted corrections.

24. What does “Apply Operation – Period Range” mean in the Period Setup dialog?

It maps each target month to its corresponding source month within the defined range. It's a positional copy — the first month of the target range gets the value from the first month of the source range, and so on.

25. If I apply a +3% growth to Beer at the sub-total level, does it apply uniformly across all Beer brands or proportionally?

Proportionally — each brand within Beer receives the growth distributed according to its share of the category total, so the existing brand mix is preserved. If it needs to be distributed based on any other measure or equally, this can be configured after entering values in the cells.

Part 4

Writeback

26. If a user isn't able to add a writeback destination from a planning sheet, is that always a writeback configuration issue, or could it be a permissions issue?

It can be either, so permissions should be checked first. Writeback-specific access is controlled separately under Security → Writeback, so a user missing from that list would be blocked from writing back regardless of how correctly the destination itself is configured.

27. Where is writeback access for a specific user actually granted?

Under Security, on the Writeback tab, users are searched for by name or email and added individually to the list of people permitted to write back data from that planning sheet.

28. Is writeback access tied to general report access, or is it a separate permission layer?

It is a separate layer. General report access controls who can open and interact with the planning sheet, while the Writeback tab under Security specifically controls who is permitted to write data back from it.

29. Does writeback access need to be configured before or after adding a writeback destination?

The destination can be added at any time, but writeback access under Security → Writeback must be granted before a user can actually run writeback. A user who can see and configure the destination but is not on the writeback access list will be blocked when they attempt to write back — the configuration is correct but the permission is missing.

30. What is the difference between Long, Wide, and Long with Changes writeback formats?

Long format writes one row per data point — each combination of dimension members and time period is a separate row. Wide format writes one row per dimension combination with each measure or time period as a separate column — closer to an unpivot table structure. Long with Changes writes only the rows that have been added or modified since the last writeback, making it faster and more storage-efficient for large datasets.

31. Why do we uncheck Gross Revenue in the Data tab before running writeback?

The Gross Revenue measure is a read-only actuals measure pulled from the semantic model — it is not plan data and does not need to be committed to the SQL database. Unchecking it ensures only the Forecast measure is written back, keeping the destination table clean and focused on planning data.

32. Can I write back to more than one destination from the same planning sheet?

Yes — a single planning sheet can have multiple Fabric SQL DB destinations configured. Each destination can have its own writeback type, column mapping, and filter settings, and can be toggled on or off independently from the Destinations tab in writeback settings.

33. Can I write back multiple sheets to the same SQL table?

Yes — multiple sheets from the same Plan app can be written back to the same SQL database. The SQL DB destination needs to be configured in every sheet that needs to be written back.

34. Can I control which columns get written back?

Yes — in the Data tab of the writeback settings, you can choose which columns to include or exclude. This is useful if you have workflow columns (Submitted, Approval Status, Comments) that you want to write back alongside the plan data, or if you want to exclude certain columns from the destination table.

35. What does Decimal Precision control, and does it affect the values in the Plan app?

It controls how many decimal places are stored in the SQL table. It does not affect the values displayed or used within the Plan app — it only impacts the precision of the written data in the database.

36. What happens if I run writeback twice — does it create duplicate rows?

It depends on the writeback type. For Long and Wide formats, running writeback again replaces the existing rows with the current values — no duplicates are created. For Long with Changes and Wide with Changes, a new row is added only when data has actually changed — if nothing has changed since the last writeback, running it again produces no new rows. This is what makes the With Changes formats useful for audit trails: each change is recorded as a new row, while unchanged data is left as is.

37. What information does the writeback log capture?

It records the status (success/failure), duration of the writeback, the user who triggered it, the measures that were written back, and the writeback type used. This provides a full audit trail for governance and troubleshooting.

38. What happens if writeback fails partway through — is any data written?

The logs will show a failed or partial status. It's best to investigate the log details, resolve the issue, and re-run writeback rather than assuming the table is complete.

39. If I verify the SQL table and the data looks wrong, what should I check first?

Start by checking the writeback logs in the Writeback ribbon — if the writeback failed or completed only partially, the logs will show the status, duration, and any errors. If the writeback completed successfully, check the writeback type — Long vs Wide affects how data is structured in the table. Also confirm you are viewing the correct scenario in the SQL table — if Scenario 1 was written back but you are previewing the Base scenario rows, the values will appear mismatched. Finally, verify the filter type was not unintentionally set to something restrictive and confirm the correct measures were selected in the Data tab before running writeback.

40. If I delete a row in the planning sheet, does writeback delete it from the SQL database?

No — writeback is additive. Deleting a row in the planning sheet does not remove it from the destination table. To remove data from the SQL database, you would need to do so directly in the database.

41. Can I write back to a destination outside of Fabric SQL — for example, to a SharePoint list or an Excel file?

No — Fabric Planning writes back exclusively to Fabric SQL databases. For downstream consumption in other tools, the data in the Fabric SQL database can be queried directly or exposed via a semantic model, Power BI report, or dataflow.