Explain the gap in an afternoon

Your forecast, budget, and pipeline each said one thing. The actuals said another. This playbook reconciles all five versions of the number and explains the gap, using prompts you run yourself.

START WITH FIVE DISTINCT DATA POINTS

Actuals
what happened

Forecast
what you expected

Budget
what you committed

Pipeline
what sales implied

Assumptions
what you bet on

The gap between them, explained line by line and assigned to an owner

Before you run anything: open any AI assistant your company allows, and decide your materiality threshold (say 10% and $50K).

THEN RUN THE FOUR PROMPTS IN ORDER

1 Run Prompt 1: the variance bridge

Paste actuals against forecast and budget; keep pipeline handy for revenue lines. This quantifies the gap line by line and labels each timing, performance, forecast error, or unresolved. Save the table, you reuse it in step 4.

PASTE IN: Actuals, forecast, budget

YOU GET
The gap between forecast, budget, and actual, quantified and classified

2 Run Prompt 2: the accuracy scorecard

Paste six months of forecasts and actuals. This tells you which lines deserve trust and which have a habit of missing.

PASTE IN: Last 6 months

YOU GET
A trust map: which forecast lines to believe, which to challenge

3 Run Prompt 3: the assumption audit

Paste the assumptions behind the forecast, including the pipeline that was meant to convert, against what happened. Grades each one held, partial, or failed.

PASTE IN: Assumptions, pipeline, actuals

YOU GET
Which bets broke, including pipeline that did not land, and the fix

4 Run Prompt 4: the review pre read

Paste the bridge from step 1 and your owner list. Assembles the pre read your meeting actually opens on, every gap assigned.

PASTE IN: Step 1 bridge + owners

YOU GET
A pre read that starts on decisions, not on slide one

WHAT YOU WALK IN WITH

Every gap between what you forecast, budgeted, and pipelined versus what actually happened, explained, quantified, and assigned. Roughly 45 minutes, start to finish.

The four prompts are on the next page. Copy, paste, and add your data where bracketed.

The four prompts

Run them in the order on page one. Paste each into your AI assistant and fill the brackets with your own data and thresholds.

Prompt 1: Variance bridge with enforced evidence

Act as an FP&A manager preparing the monthly variance review. Input: actuals versus forecast by line item: [paste data]. Materiality: explain lines beyond [X percent and $Y]. Build a bridge from forecast to actuals ordered by absolute dollar impact. Classify each material variance using strict rules: "timing" requires an identifiable reversal, name the period it reverses; "performance" requires a business driver visible in the data, name and quantify it; "forecast error" requires identifying which assumption was wrong and by how much. Any variance not meeting one of those evidence bars is classified "unresolved" with the specific question that would resolve it. Output a table with columns: line, variance, classification, evidence, quantified driver, reversal period or owner question. Variances must sum to the total; show the tie out.

Prompt 2: Forecast accuracy scorecard

Act as an FP&A director assessing whether the forecast process deserves the confidence placed in it. Input: the last six monthly forecasts and corresponding actuals, by line item: [paste data]. For each line compute mean absolute percentage error, bias (average signed error, stating whether the line is systematically forecast high or low), and hit rate within a [tolerance] percent band. Show the formulas once, then apply throughout. Identify the three worst lines and describe each error pattern in one line: chronic sandbagging, chronic optimism, or unpredictable volatility. For each problem line recommend one process change: a different driver, a different owner, or a wider acknowledged range. Present as a table plus a three sentence summary a CFO would read as: here is how much to trust this forecast and where.

Prompt 3: Assumption audit

Act as an FP&A manager closing the loop between assumptions and outcomes, the step most forecast processes skip. Input: the stated assumptions behind the current forecast and the actuals now observable: [paste both]. Grade every assumption held, partially held, or failed, citing the specific actual justifying each grade. For failed assumptions, quantify the forecast impact using the data provided and show the arithmetic. Assumptions with no observable actual yet are listed separately as untested, not graded. For each failed assumption propose the replacement for next cycle in one sentence with its rationale, and state the confidence basis: grounded in observed data, or judgment pending more evidence. Close with the one assumption whose failure would hurt most next quarter and what early indicator would reveal it failing.

Prompt 4: Review meeting pre read

Act as the FP&A lead writing the pre read for the monthly forecast review, audience: CFO and budget owners. Input: the variance bridge or table: [paste data], and the budget owner list: [paste owners]. Structure, under 600 words: a three sentence summary of the month a CFO could repeat verbatim to the CEO; the five variances that matter ranked by dollar impact, one short paragraph each, every explanation citing a figure; a decisions needed section listing the two or three choices this meeting must actually make, each framed as a question with options; and open questions assigned to named owners with the date answers are due. Anything unexplained goes in open questions, never in the narrative. The test of a good pre read: the meeting starts at the decisions section, not at slide one.

Tip: keep each output in one document. Step 4 reads best when it can see the bridge from step 1.

When doing it by hand gets old

The playbook works. It also resets to zero every month. Here is what changes when it does not.

YOU DID THE HARD PART

Running these four prompts once proves the outcome is real: the variance review, explained and decision ready, in an afternoon. The catch is that next month you export everything again, paste it again, and check the columns again. The thinking is solved. The plumbing is not.

Connected, not pasted

The same analysis runs against your ERP, planning, and CRM directly, so there is nothing to export.

On a schedule

The pre read lands in your inbox before the review, every month, without anyone kicking it off.

Traceable

Every number links back to its source, so the answer survives a controller and an auditor.

A 30 MINUTE CALL

Walk us through your variance review.
We will map a clear path to running it for you.

Book the call at sales@nextq.ai. Bring your use case; leave with a plan, not a pitch.

Keep this playbook. It works with or without us. The 30 minute call is open whenever you want to map the path.