

The Evidence Chain

AI READINESS · COACH IT INTO NUMBERS

Your team's AI access is a budget line somebody else controls. Whether you are defending the access you have or building the case to expand it, the decision turns on numbers, and enthusiasm does not guarantee resources. Follow this working session fixes that.

A renewal date, a budget review, or a license or token expansion request is the signal. The leader who converts champion enthusiasm into system-verified numbers before that conversation keeps the tools, grows the seats, and holds the credibility either way.

COMPANION ARTICLE · The full walkthrough, good and bad output examples, and the systems glossary live at schumacherconsultancy.com. This page is the working session itself.

STEP 1 · DISCOVER · PASTE INTO YOUR GENAI/LLM (LARGE LANGUAGE MODEL) APPLICATION

ROLE

You are a higher education operations analyst who designs measurement for change initiatives. You treat memory and self-report as inadmissible evidence.

THINK FIRST

Do not generate yet. Interview me one question at a time, minimum 5 questions, maximum 8, until I can name my 3 heaviest AI users and what they produce, the systems my unit already runs, and my license decision date. Then restate my situation in exactly 3 sentences and wait for my confirmation.

CONTEXT

Team of [#]. ERP/CRM/SaaS: [tool]. Deployed [# months] ago. Decision date (renewal, budget review, or expansion request): [date].

TASK

After I confirm, identify 3 to 5 recurring deliverables where a before-and-after number already sits in a system we run.

CONSTRAINTS

Reject any candidate whose number would come from memory or self-report. A metric without a system timestamp dies at the first budget-officer question. Mark every owner you have not confirmed with me as UNVERIFIED. Output your findings as a clean markdown table I can highlight, copy, and paste into Excel or Google Sheets with each value landing in its own cell, with these five columns in this order: Deliverable, Metric to verify, System of Record, Owner, Keep (Y/N)? Leave every Keep (Y/N)? cell blank. Output only the table. No text before or after it.

THINKING PAUSE

In a <thinking> block first, list your 3 most fragile assumptions about my unit and name the one deliverable most likely to collapse under scrutiny.

RECEIPT

Close with a yes/no audit: interviewed before generating / every deliverable names a system / every unconfirmed owner flagged / nothing outside the context block.

STEP 2 · THE HUMAN CHECKPOINT

The LLM generated a table. Bring it into Excel or Google Sheets and verify it in place. There is no grid to retype. Nothing UNVERIFIED makes it to Step 3.

1. In your LLM, drag to highlight the whole table, then press Cmd+C.

2. Open a blank Excel workbook, click cell A1, and press Cmd+V. Each value lands in its own cell. If everything lands in one column instead, undo (Cmd+Z), then paste with Edit > Paste Special > Text.
3. In the Keep column, type Y for any number you located in its named system within 5 minutes. Type N for the rest.
4. Right-click each N row number and choose Delete. Only verified rows move to Step 3.
5. Save the workbook to your team's Teams channel or SharePoint site. This file becomes your verification record.

THREE WAYS YOUR DATA FAILS: Vanity metrics (impressive, changes nothing) · Self-report (a feeling with a number attached) · Undefined owner (a role, a team, or anyone who is not a name)

STEP 3 · GENERATE · PASTE INTO SAME LLM CHAT

No OKR practice in your unit? Read Objective as goal and Key Result as evidence. The prompt works identically.

ROLE

Same operations analyst. Memory and self-report remain inadmissible.

CONTEXT

Below is a table of deliverables I have verified, pasted from a spreadsheet. It may arrive as plain rows separated by tabs or spaces: treat the first line as headers and each line after it as one row. Every row is marked Keep = Y, meaning I confirmed the metric, system, and owner myself. If any row is marked N or has a blank Keep cell, ignore it completely and do not fill in or guess at anything from it.

[PASTE YOUR UPDATED TABLE HERE]

TASK

Open with one line: who this OKR set defends, the measurable outcome by [decision date], and the single strongest proof point in the block. Then write up to 3 objectives describing how the work changes, each with up to 3 key results reportable by [decision date]. After drafting, ask me one question at a time for the target of each key result: the number we commit to reach by [decision date], plus my basis for it in one phrase (trend, benchmark, or mandate).

CONSTRAINTS

Every key result names its baseline, source system, and owner from the block. Audit every key result for undefined scope qualifiers (eligible, routine, comparable, standard, key): replace each with the explicit system field or category codes that define it, or append exactly one flag: DEFINE: [term]. The owner must answer: which items count, by what field? A key result missing baseline, source, or owner is labeled NOT YET MEASURABLE. Never invent a target. Targets come from me. A key result I have not given a target is labeled NO TARGET SET and cannot appear on the dashboard.

TRACKING RECORD

After I supply targets, build my tracking record. First compute the checkpoint schedule from today and [decision date]: under 4 weeks of runway, checkpoints are weekly. 4 to 12 weeks, every 2 weeks. Over 12 weeks, monthly. Always include a final pull inside the 5 business days before [decision date]. Print real calendar dates. Never print Day 30, Day 60, or any placeholder interval. Output one block per key result. Each block opens with exactly these lines: KEY RESULT. MEASURING (one plain sentence naming what this number tells a VP). UNIT (hours, days, count, percent, or dollars). SOURCE SYSTEM. OWNER. HOW TO PULL IT (the exact report, field, and filter the owner opens). CALCULATION (direct read, or the formula). TARGET (number, unit, decision date, and my basis). Under each header, one clean markdown table I can highlight and copy. Column 1 is the pull label. Column 2 is the date. Then one column per raw input the calculation needs, each labeled with what that input is and its unit. Then one column for the metric itself, labeled with its name and unit. Then one column, PULLED BY. A direct-read metric has one input column. A derived metric has one column per input, and the metric column holds the computed result, never a raw input. Row 1 is Baseline, carrying the before number from the verified block and the date that number describes. Every row after it is one computed checkpoint date. Leave every cell blank except the Baseline row and the dates.

THINKING PAUSE

In a <thinking> block first, name the key result most likely to be gamed and exactly how someone would game it.

RECEIPT

Close with a yes/no audit: only verified context used / no undefined qualifier left unflagged / no invented targets / every key result carries baseline plus source plus owner.

STEP 4 · THE COMMON SENSE CHECK

The LLM audited itself. Now you audit your LLM. Three questions, yes or no. A no on any question sends that key result back, not forward.

1. Would you say this number out loud in the review meeting?
2. Could the named owner locate it in the system within 5 minutes?
3. Does it describe a change in (not effort spent on) the work?

STEP 5 · MAKE IT PERMANENT

Save all three prompts with your verified context filled in: that is your unit's measurement intake, reusable for every future tool.

Step 3 is your tracking record, one block per key result, checkpoint dates computed from your decision date. That record is the only data the dashboard reads. Save it as a file in your team's cloud (i.e., Teams/Slack channel, OneDrive, Google Drive), not your personal drive. Leaving this on your laptop and not the cloud just reserves extra time and steps for you downstream.

Now anchor it. Send every checkpoint date in your record as a calendar invite to the owner named on that block, with a link to the shared file in the invite body. Paste that block's HOW TO PULL IT line into the invite body so the owner never has to ask what you meant. The owner opens the invite, follows the line, pulls the number, enters it in the row, and initials PULLED BY. An OKR set without calendar checkpoints is a document, not a dashboard.

STEP 6 · DASHBOARD · PASTE INTO YOUR LLM TOOL AT EACH CHECKPOINT

ROLE

You are a data storyteller who builds executive dashboards on Storytelling with Data principles. Your audience is a VP with 30 seconds.

THINK FIRST

Do not generate yet. If any key result in my record is missing a baseline, target, or at least one checkpoint value, ask me for it first. A chart cannot show progress toward a target that does not exist.

CONTEXT

My tracking record is below, one block per key result, each with a header and a measurement log. Decision date: [date]. Today: [date].

[PASTE YOUR UPDATED TRACKING RECORD HERE]

TASK

Produce one complete, self-contained HTML file showing each key result's progress from baseline toward target across the checkpoints recorded so far. Open the page with one summary line: how many key results are on pace, how many are behind, and days remaining to the decision date.

CONSTRAINTS

Output only the HTML code in a single code block, nothing before or after it. Inline CSS and SVG only, no external libraries, no internet required, so the file opens in any browser. Choose each chart by the data. Chart the metric column only. Never chart a raw input column. A key result with a baseline plus one or more filled checkpoint rows is a line from baseline through the latest value, with the target as a labeled horizontal reference line. A key result with a baseline and no filled checkpoints yet is a horizontal bar from baseline to target with no current position marked. Take every axis label and every printed unit from that block's UNIT line. Write every takeaway headline from that block's MEASURING line, in plain words. Never use pie charts, donuts, gauges, or 3D. One accent color for progress, gray for everything else. If a key result is behind the pace needed to reach its target by the decision date, switch its accent to a warning color and say so in its headline. Never a label like KR1 Progress. Print baseline, current, and target as numbers beside every chart. Label data directly on the chart. No legends, no gridlines, no chart borders. List any key result whose rows are all blank below Baseline at the bottom of the page under one heading: Awaiting first pull. Close the page with one footer line: Built with the Evidence Chain · schumacherconsultancy.com.

THINKING PAUSE

In a <thinking> block first, check each key result's trajectory against its target and the decision date, and name the one chart a VP will look at first.

RECEIPT

Close with a yes/no audit: every chart has a takeaway headline / target visible on every chart / one accent color plus gray only / no pies, gauges, or legends / single self-contained HTML file / nothing output except the code block.

SAVE AND SHARE THE DASHBOARD

1. Your LLM returns a code block. Click its copy icon.
2. On Mac: open TextEdit, press Cmd+Shift+T for plain text, paste, and save the file as Dashboard.html. On Windows: open Notepad, paste, and save as Dashboard.html.
3. Drop the file in the same Teams/Slack channel or SharePoint/Google Drive folder as your tracking record. Anyone who opens it sees the current picture. Regenerate after every checkpoint so the story stays current.