


THOUGHT LEADERSHIP

See Strategy Execution Drift Before Leadership Does

Strategy loses value in the weeks between something changing and the right person finding out. Harmony, the AI layer in Conductor, is built to shorten that time.

Al Torreno September 17, 2026



Strategy loses value in the weeks between something changing and the right person finding out. Harmony, the AI layer in Conductor, is built to shorten that time.

THE STRATEGY EXECUTION GAP

Where strategy loses its value

It's the fourth portfolio review of the year. An initiative that was green for three consecutive reviews has just gone amber, and the room wants to know what happened this month.

It started weeks ago. A milestone moved. A dependency changed in another workstream. Finance revised the benefit forecast in a separate spreadsheet, and the program team didn't see it until this cycle. Everyone is making reasonable decisions. They're just making them from different versions of reality.

Multiply that room across a portfolio and the cost adds up.

Strategy rarely fails because it was wrong. It fails in execution. Only 63% of a strategy's intended benefits are realized. The other 37% is lost to failures in coordination and information.

EXHIBIT 1

The cost of the status quo

63%

of a strategy's intended benefits are realized.

MANKINS & STEELE, HARVARD BUSINESS REVIEW, 2005

37%

is lost to failures in coordination and information.


MANKINS & STEELE, HARVARD BUSINESS REVIEW, 2005

Traditionally, strategic plans get reinterpreted at every layer they pass through, and by the time they reach the people doing the work, the original intent is unrecognizable. A date moves, scope gets trimmed, a target gets restated in someone else's words, and the original plan quietly becomes a different plan without the governance that once surrounded it. When the plan lives in slides and spreadsheets, drift surfaces only when somebody assembles the next review pack (reactive).

That is execution drift. It lives in the time between a change happening and the moment someone who can act on it finds out.

EXHIBIT 2

With Harmony, drift gets caught while there is still time to correct it



CONDUCTOR BY SENSEI LABS

What would it take to catch that slip at the first review instead of the fourth?

Harmony is the AI layer in Conductor. It understands what your organization has recorded, compares it against what great strategy execution looks like, and identifies where the plan and reality have started to drift. It tells the right person, in the right order, gives them the context and recommendations they need to make a decision, and, once approved, moves that decision forward.

Closing the strategy-to-execution gap guides our design of Harmony AI. Every capability described here, including reporting and visualization, shortens the distance between strategic intent and the reality of execution.


It starts with the system knowing what the initiative is connected to.

THE FOUNDATION

How Harmony knows what an initiative is connected to

Harmony can only shorten the time to action if it understands what changed, what that change affects, who owns it, and what should happen next. That's why the foundation comes first.

For our initiative, that means Conductor knows its milestones, the dependency sitting in another workstream, the benefit forecast finance owns, and the person accountable for each.



A model of your business. Conductor has a standard data model: projects, gates, KPIs, risks, etc. come out of the box with standard definitions. Custom objects let you extend the standard model with the items unique to your business (e.g. business units, factories, customers, invoices, etc.) Conductor ends up as a digital twin of your strategy, governance model, business objectives, strategy execution logic, KPIs, and project health and progress. That structure is what turns a question like “which initiatives in my workstream are waiting on a regulatory dependency, and who owns each one” into something the system can answer without anyone having premeditated a dashboard.



Your data. A single source of truth is even more powerful when you increase the number of sources feeding it, and the things that make a program slip often live in systems that Conductor cannot see. The HRIS. The finance system. The spreadsheet a workstream lead keeps on their own laptop. The meeting where a date moved and nobody updated a record. The integration framework is how it's possible, and choosing which systems Conductor can connect to is up to you.

Knowledge. Harmony can only judge the work against a standard, and the standard has to be yours. Conductor ships with knowledge on what strategy execution best practice looks like. You layer on top. Everything you configure carries its own description: a stage gate explains the purpose of its review and the decision standard, a custom field explains what belongs in it. The playbooks, operating standards, process guides and policies that define what good looks like get authored in Conductor, and Harmony reasons over them the same way it reads the descriptions.

Tools. Calculated fields, the rules that read them, reports and scheduled jobs. Arithmetic, thresholds and pass or fail live here rather than in a language model. Conductor also exposes its data and actions through MCP, the standard interface AI assistants use to reach other systems, so your people can work with Conductor from whatever tool they already have open. The MCP respects the user's access and permissions in Conductor.

Language and reasoning. Plain language in, and an answer shaped to the question. A grid, a chart, a ranked list, a draft, or an explanation of why something is being flagged.

The insight, recommendation, and decision. Drift gets caught while there is still time to correct it. A workstream lead learns what a gate actually expects before the review rather than during it. A sponsor stops second-guessing a green rating because the evidence arrived attached to it. Trust in the data grows the way it grows anywhere, by being right when it mattered. Harmony puts that moment wherever the decision actually gets made. The homepage inbox beside your approvals. An initiative page. A dashboard widget. A scheduled digest. An answer in Harmony chat.

Trust in the data grows the way it grows anywhere, by being right when it mattered.

Al Torreno

VP Solution Delivery, Sensei Labs

WHAT HARMONY DOES

Five jobs that shorten the time to action

Five jobs, described the way they'll show up in your program. Here they are in the order they would happen for our initiative.


01

Recommend

The week the dependency moves, the initiative lands on the sponsor's list.

The five initiatives most worth worrying about this week, ranked, with the reasoning shown. A green rating next to a cautionary review comment and a missed target is the kind of contradiction Harmony is good at catching.

In practice: An initiative that was green for three consecutive reviews and has just gone amber. Monthly benefit values that no longer add up to the annualized figure on the business case.



02

Explain

The owner needs to know what the next gate will expect of a recovery plan.

Ask how a gate works, what a Gate 2 review examines, or which fields are required, in your own words. Answers are grounded in your guidance and cite it. Where the guidance does not exist, Harmony says so instead of inventing an answer.

In practice: An initiative owner asks what Gate 3 actually examines and gets the gate's own review scope back, cited. A workstream lead asks why an initiative is not yet eligible to submit for Gate 4 and gets the specific conditions that are unmet.

03

Draft

The owner opens the review with a first version already in place.

Turn rough bullets into an initiative description. Propose titles and likely stakeholders. Write the first version of a gate narrative from data already sitting in the record. The author edits and owns it.

In practice: The business case arrives with a first pass already written. Baseline, assumptions, risks and readiness drafted from what the record holds, so the owner opens something to sharpen instead of an empty tab. A review comes with its rating recommended from KPI variance, schedule slip, gate age and open risks, and the commentary already drafted underneath it.

04

Interrogate

The sponsor wants to know what else is sitting on that same milestone.

Ask a question in plain language. Harmony traverses the relationships, builds the query, and returns a grid, a chart, or a written answer depending on what suits the question.

In practice: Ask for CAPEX variance by workstream across initiatives past Gate 3. Harmony writes the query, chooses the visualization, and renders it with the same charting the reporting portal uses.

05

Configure

Next time, nobody has to flag it by hand.

Describe what you want to see and get a working dashboard or alert, drawing on patterns built for comparable programs elsewhere.

In practice: Describe a rule in plain language, such as the milestone conditions to satisfy before Gate 2 submission, and get the configuration back as a change you accept or reject. Update a value across many initiatives from chat in one confirmed pass rather than record by record.

INSIGHT IN PRACTICE

The answer shows up where the decision gets made

All of this reaches the sponsor where they already work.


Everything above needs somewhere to display. Reporting, dashboards, and visualization live natively inside Conductor, not in a separate BI tool fed by exports (although you have that option as well).

The reporting comes from the work itself, so nobody builds six dashboards with six storylines.

When the sponsor wants to see the pattern across the whole portfolio, they ask for it.

Harmony can help you write reporting portal reports: turn a plain-language request into draft SQL, or get an existing report explained back to you in plain language. This is Draft and Interrogate applied to reporting. The data you need moves from days or weeks to minutes or hours.

This is Configure applied to visualization: describe what you want (“actual spend by portfolio, each quarter”) and Harmony proposes the data source and chart together, with server-side rules keeping the pairing sensible so an unsuitable chart is never built.




A filter set at the dashboard level follows the viewer across every dashboard page until changed; a filter set on a single dashboard page narrows just that view. Either way, filter values flow into the underlying reporting portal reports automatically. A viewer picks a portfolio once, and every widget fed by that filter re-runs against it, with no reconfiguration.

GUARDRAILED BY DESIGN


What stays under your control

None of that required anyone to see something they aren't cleared to see.

Harmony does the analysis on behalf of the person and respects that person's permissions and access. This will ensure insights are relevant. This will also enforce the security and confidentiality of data.



Your data never passes through the model itself. Harmony writes the recipe, and Conductor's own engine runs the query and draws the chart.



Reports are guardrailed, so your data stays safe, and every report runs in the context of the viewer, so nobody sees a project, portfolio, or row they aren't authorized to see.

FOR THE EPMO

What changes for the EPMO

EXHIBIT 10

The same review cycle, two ways

TODAY

- Someone manually chases every update before the review
- Drift surfaces during the next review cycle
- A workstream lead learns what a gate expects during the review
- A sponsor questions a green rating because the evidence is scattered
- A new question becomes a reporting request
- The portfolio view waits for the next review pack

WITH HARMONY

- The update comes from the work itself
- Drift is flagged while there is still time to act
- They know the requirement before submitting
- The evidence sits with the rating
- The sponsor asks and gets the answer
- The sponsor can ask the portfolio a question at any point in the cycle

THE PAYOFF

Back in the review room

A few months later, another initiative slips, and this time the room already knows. Harmony flagged the moved dependency three weeks ago. The owner brought a recovery plan, and the sponsor has read the evidence behind the amber. The conversation is about what to do next.

This is the strategy-to-execution gap closing in practice: a sponsor steers a dashboard instead of waiting for the next review pack, and the answer that used to take a manually assembled deck now takes a filter click and, increasingly, a plain-language question to Harmony.

Most tools tell you what happened. Conductor changes what happens next.

Harmony's capabilities are rolling out across Conductor Modern in stages.

Book a Demo

See how Conductor gives leaders full control over complex transformations

Schedule Now

Frequently asked questions

What is Harmony AI?

Harmony is the AI layer in Conductor. It understands what your organization has recorded, compares it against what great strategy execution looks like, and identifies where the plan and reality have started to drift.

Does my data pass through the AI model?

Your data never passes through the model itself. Harmony writes the recipe, and Conductor's own engine runs the query and draws the chart.

Can Harmony show someone data they aren't allowed to see?

Harmony does the analysis on behalf of the person and respects that person's permissions and access. The MCP respects the user's access and permissions in Conductor.

What happens when Harmony doesn't know the answer?

Where the guidance does not exist, Harmony says so instead of inventing an answer.

Does Harmony make changes on its own?

No. Configuration comes back as a change you accept or reject, and updates across many initiatives run in one confirmed pass.

READ ONLINE

Read this article on our website

getconductor.com/resources/see-execution-drift-before-leadership-does



GET STARTED

Talk to sales

Walk through an example with our team and see Conductor on your own programs.

[Talk to Sales →](#)



Scan to book