THE 2026

Signal Architecture *Framework.*

How the Google ecosystem, Meridian-grade measurement, and agentic AI fit together as one system — and the audit, readiness, and decision-economics tools to operationalize it.

By Kyle Schwietz · Founder, Uncommon Move

uncommonmove.com

CONTENTS

What you get.

Three connected tools, designed to be used inside a real organization, by a real team, in a real budget cycle. Each one is the connective tissue between conversations that usually happen in separate rooms.

The arrangement is intentional. Each tool builds on the one before it. Read in order; revisit out of order.

HOW TO USE THIS DOCUMENT

Three tools. One *system.*

This Framework exists because the conversations about Google ads, marketing measurement, and agentic AI have been happening in three separate rooms — and the people who win the next decade are the ones who put them in the same room. It is not a comprehensive guide to any one of those topics. It is the connective tissue between them: the architecture, the audit, and the decision economics that make the three work as one.

Each of the three tools below is designed to be used inside a real organization, by a real team, in a real budget cycle. They are deliberately short. They are deliberately opinionated. They are deliberately translatable into language the rest of the C-suite uses.

Tool 1 · The Signal Flow Map

How signals actually move through the Google ecosystem — YouTube and Search feeding Demand Gen and PMax, with Meridian and GA4 underneath as the proof layer. Use this to brief leadership on why the channels are not the strategy — the flow between them is.

Tool 2 · The Meridian Readiness Checklist

Twelve checks that determine whether your organization is ready to run Marketing Mix Modeling that will actually be acted on. Most teams pass the technical checks and fail the organizational ones. Use this before the Meridian engagement, not after.

Tool 3 · The Cost Per Decision Worksheet

The CFO-readable version of the agentic AI conversation. A four-input model that translates marketing decision velocity into financial-engineering language — Cost Per Decision, equivalent FTE capacity, share of marketing payroll. Use this to justify the AI investment to people who do not buy productivity arguments.

TOOL 1 OF 3

The Signal Flow Map.

Most teams treat Google's products as separate channels with separate budgets and separate teams. That is the silo problem, and it costs compounding performance nobody sees.

The Google advertising ecosystem is designed to work together. YouTube signals feed Search. Demand Gen feeds the consideration layer. Performance Max synthesizes everything. And Meridian MMM is the layer that proves the value of the whole — so the organization can invest with confidence rather than guesswork.

How to read the map

Four tiers, top to bottom: **Signal Sources** (YouTube, Search/Shopping, DV360/CTV) generate the raw signal. **Activation** (Demand Gen, Display) shapes that signal into mid-funnel and reach campaigns. **AI Synthesis** (Performance Max) is the layer Google's AI uses to run across all inventory simultaneously. **Proof Layer** (Meridian MMM, GA4, incrementality, attribution) measures whether any of it actually worked.

The signal flow rule

Every dollar of media performance is a function of the signal you feed the system. The campaigns are the symptom. The signal architecture is the cause. **Build for the signal flow first; the campaign performance follows.**

Where most organizations break the map

Break 1 — Channel-siloed teams. Search is owned by one team, YouTube by another, Display by a third. No one is responsible for the flow between them. Result: PMax is fed inconsistent signal because no one is engineering the pipeline that feeds it.

Break 2 — Treating PMax as a campaign. PMax is not a campaign — it is the synthesis layer. Treating it like another tactic with creative refreshes and bid adjustments misses the actual mechanism. PMax compounds on signal quality, not campaign management.

Break 3 — No proof layer. Without Meridian or rigorous incrementality testing underneath the stack, the organization cannot defend the investment in the upper-funnel signals (YouTube, DV360) that are doing the actual work. Budgets get cut from the channels that matter most because they are the hardest to attribute.

TOOL 1 · CONTINUED

The map, as a one-page brief.

Use the diagram below to brief executive leadership. The point is not the visual — it is that they leave the meeting understanding that the channels are not the strategy. The flow between them is.

SIGNAL SOURCES

YouTube · Search & Shopping · DV360 & CTV

Watch signal · intent signal · reach signal

↓

ACTIVATION

Demand Gen · Display

Mid-funnel synthesis · programmatic reach

↓

AI SYNTHESIS

Performance Max

Google's AI running across all inventory simultaneously

↓

PROOF LAYER

Meridian MMM · GA4 · Incrementality · Attribution

The infrastructure that proves the value of the whole stack

The arrows are the strategy. Most marketing teams own the boxes. Almost no one owns the arrows.

Three diagnostic questions.

The map is useful only if it leads to a different conversation. Use these three questions to find out where your organization actually sits on the flow.

1. Who owns the signal pipeline?

Not the channels. The pipeline. If the answer is "everyone" or "no one," that is the first thing to fix. Without a single owner of the signal architecture, PMax will always underperform its potential because no one is engineering what it sees.

2. What is your proof layer cadence?

How often do you refresh your MMM? How often do you run incrementality holdouts? How often do you audit GA4 for the structural biases in the next section of this Framework? If the answer is "annually" or "when something looks wrong," the proof layer is decorative, not operational.

3. Where does the signal flow break first under stress?

Run the thought experiment: budget gets cut 20%. Which channel gets cut first? The one that is hardest to attribute, almost certainly. Now run it again with Meridian results in hand. The answer should be different. If it is not, the proof layer is not changing decisions — and the rest of the architecture cannot do its job.

TOOL 2 OF 3

The Meridian Readiness Checklist.

Twelve checks. The technical ones are obvious. The organizational ones are why most Meridian engagements produce a model nobody acts on.

Marketing Mix Modeling is the measurement discipline most enterprise organizations skip — not because it is optional, but because implementing it well requires data infrastructure, organizational alignment, and leadership willingness to act on findings that may challenge politically settled budget allocations.

Meridian, Google's open-source MMM framework, has changed the technical side of that equation: the model is transparent, the priors are inspectable, the refresh cycle is measured in weeks rather than quarters. But the organizational side has not changed. **Run this checklist before you commission the model — not after.**

How to score this checklist

Each item is a yes / no / partial. Score 12/12 and you are ready to commission a Meridian engagement that will produce findings the organization will actually act on. Score under 8/12 and you are not — the model will be built, the findings will be presented, and the budgets will not move. Fix the gaps first.

TOOL 2 · TECHNICAL READINESS

The six technical checks.

01. Time-series spend data, by channel, at consistent granularity

At least two years of weekly (or daily) spend data, broken out by channel — not by campaign, not by platform, by channel. Inconsistent granularity (daily for some, monthly for others) breaks the model.

02. Clean outcome data tied to the same time series

Revenue, conversions, leads, or whatever business metric you are modeling — at the same cadence as the spend data, with the same timezone, with the same definition across the entire window. Definition drift ("we changed how we count a lead in Q3 2024") corrupts the model silently.

03. External variables documented and accessible

Promotions, price changes, competitor events, macroeconomic indicators, seasonality, weather (if relevant). Without these, the model attributes their effects to your media — and you will believe your campaigns are driving things that the calendar is driving.

04. Reach and frequency data from Google's first-party signals

Meridian's advantage over legacy MMM tooling is the ability to incorporate Google's own reach and frequency data. If you have not connected your Google Ads, DV360, and YouTube accounts to the modeling environment, you are giving up the model's biggest edge.

05. Internal Python / Bayesian capability — or a credible partner

Meridian runs in Python. Inspecting the priors and adjusting the model requires technical fluency. Either your data team has it, or your engagement partner has it and is willing to be transparent about the model. Black-box partners defeat the entire point of open-source MMM.

06. A documented refresh cadence

Monthly, quarterly, or event-driven. The cadence matters less than the fact that one exists and is on someone's calendar. A model that is built once and never refreshed becomes wrong faster than anyone expects.

TOOL 2 · ORGANIZATIONAL READINESS

The six organizational checks.

These are the ones nobody warns you about. Failing any of these means the model will get built, the deck will get presented, and the budget will not move.

07. Executive sponsorship at the budget-owner level

Not the analyst's manager. The person who owns the marketing budget across channels — typically a CMO, VP of Performance, or Head of Growth. Without their buy-in before the model runs, findings that challenge current allocation will hit organizational resistance and die there.

08. Pre-committed willingness to reallocate

A written or verbal commitment from leadership that, if the model finds material misallocation, the budget will move within one or two quarters. Without this, the model becomes an interesting analysis with no operational consequence.

09. Channel team alignment on the methodology

The Search team, the YouTube team, and the Display team all understand and accept that MMM measures incrementality, not platform-attributed conversions. If they have not internalized this distinction, they will dismiss findings that show their channel is less incremental than its attribution suggests.

10. Finance is in the room

The conversation about media incrementality is fundamentally a conversation about capital allocation. Finance partners who understand the methodology — and who can translate findings into balance-sheet language — are the difference between "interesting marketing data" and "board-relevant insight."

11. A clear test-and-learn protocol for validating findings

Meridian estimates incrementality. The way you confirm those estimates is controlled holdout testing. Without a documented protocol for running validation tests on the model's biggest findings, the organization will fairly question whether the model can be trusted.

12. An owner of the proof layer

A single named person responsible for the measurement infrastructure as a whole — Meridian, GA4, incrementality testing, attribution. If proof-layer ownership is split across teams or undefined, the layers will drift apart and contradict each other.

TOOL 3 OF 3

The Cost Per Decision Worksheet.

The CFO-readable version of the agentic AI conversation.

Most marketing AI conversations chase Cost Per Lead. The real lever is Cost Per Decision — the labor, latency, and capital required for an organization to reach and act on a marketing decision. Agentic AI collapses that cost by orders of magnitude, and the org structure has to be re-engineered around that economics. **This is financial engineering, not HR.**

> **Why this metric matters**
>
> Cost Per Lead lives inside the marketing department's P&L. Cost Per Decision lives inside the entire operating expense line. Reframing the AI conversation in CPD terms moves the question from "is this a marketing tool?" to "is this a structural margin lever?" — and changes who in the C-suite is willing to fund it.

The four inputs

Marketing team size (people). The headcount of the marketing organization, including embedded analysts, creative, ops, and channel managers. Used to derive how many people typically touch a single decision.

Marketing decisions per week. Campaign edits, budget shifts, creative approvals, audience tweaks, bid changes — anything someone has to decide before it happens. For most enterprise marketing orgs this number is between 100 and 400 per week, far higher than leadership intuits.

Average decision cycle time (hours). Time elapsed from "we should change this" to "the change is live." Includes meetings, analysis, approvals, and execution. The honest number is almost always longer than the team estimates.

Blended fully-loaded labor cost ($/hr). Salary + benefits + overhead, divided by working hours. For enterprise marketing orgs this is typically $100–$140/hr. Use $120/hr as a reasonable midpoint if you do not have a precise number.

TOOL 3 · THE FORMULAS

Four formulas. One spreadsheet.

Set up a spreadsheet with the four input cells from the previous page. Then add the four derived cells below. The result is the most honest single-page view of your marketing operating economics that most teams have ever seen.

Cell	Formula	What it tells you
People per decision	MAX(2, MIN(6, ROUND(1 + team / 25)))	How many people typically touch a single decision. Scales with t
Current Cost Per Decision	hourly × cycle_hours × people_per_decision	What it costs your org, fully loaded, to make one marketing decis
Agentic Cost Per Decision	hourly × (cycle × 0.30) × MAX(1, people × 0.50)	What it would cost with agentic systems handling routine work —
Annual capacity unlocked	(current_CPD – agentic_CPD) × decisions × 50	Annualized financial value of the cycle-time and human-touch re
Equivalent FTE capacity	annual_unlocked / $200,000	Translates the savings into headcount language. $200k = fully lo
% of marketing payroll	annual_unlocked / (team × $200,000)	Translates the savings into percentage-of-payroll language the C

The output is order-of-magnitude — directional, not boardroom. For a real engagement, replace the assumed 70% / 50% reductions with measured baselines from a pilot. The point of this worksheet is not the precise number; it is the realization that **the metric you have been measuring (CPL) is not the metric AI moves.**

TOOL 3 · WORKED EXAMPLE

A 40-person marketing org.

To make the math concrete, here is a typical mid-market enterprise marketing organization run through the Cost Per Decision Worksheet at default assumptions.

Input	Value
Marketing team size	40 people
Marketing decisions per week	120
Average decision cycle time	16 hours
Blended fully-loaded labor cost	$120 / hr

The math, line by line

People per decision = MAX(2, MIN(6, ROUND(1 + 40/25))) = **3**

Current Cost Per Decision = $120 × 16 × 3 = **$5,760**

Agentic Cost Per Decision = $120 × (16 × 0.30) × MAX(1, 3 × 0.50) = $120 × 4.8 × 1.5 = **$864**

Annual capacity unlocked = ($5,760 – $864) × 120 × 50 = **$29.4M**

Equivalent FTE capacity = $29.4M / $200k = **~147 people**

Share of marketing payroll = $29.4M / (40 × $200k) = **~367%**

What this number actually says

A 40-person marketing organization is not going to find $29M of cash sitting in a drawer. The number is the **capacity equivalent** — the financial value of the human time agentic AI returns to the organization. That capacity becomes either margin (if the team stays the same size and does more) or growth velocity (if the team redeploys against bigger ambitions). Either way, it is a structural lever, not a productivity tweak. **That is the conversation to have with the CFO.**

FURTHER READING

The long-form thinking *behind each tool.*

These tools are condensed. The reasoning behind them lives in the long-form essays at uncommonmove.com/thinking. If a particular tool resonated, the matching essay is the next read.

Behind Tool 1 · The Signal Flow Map

"Full-Funnel Marketing Within Google: What It Actually Means to Run the Whole System" — the long-form case for why the channels are not the strategy and the flow between them is. Reads in 8 minutes.

Behind Tool 2 · The Meridian Readiness Checklist

"What Meridian MMM Actually Is and Why It Changes How You Invest" — what Marketing Mix Modeling actually measures, why most organizations fail to act on the findings, and what the open-source release of Meridian changes. Pair this with the GA4 audit teardown for the full proof-layer picture.

"Your GA4 Setup Is Quietly Lying to You — A Teardown" — the five structural biases in default GA4 that every Meridian engagement bumps into. Required reading before checklist item 03.

Behind Tool 3 · The Cost Per Decision Worksheet

"What Agentic AI Actually Means for Marketing Teams in 2026" — beyond prompts and chatbots. How agentic systems change the unit economics of decision-making, and why that is a financial-engineering conversation rather than a productivity one.

"Stop Building for One LLM. Build for Agentic." — the infrastructure decisions that determine whether your agentic AI investment compounds or has to be rebuilt every 18 months.

Where to find these

Every essay above lives at uncommonmove.com/thinking. The Stay Sharp newsletter (uncommonmove.com) is where new long-form pieces appear first.

CLOSING

Three tools. *One uncommon move.*

If you have read this far, you are not the median marketing reader. You are the operator who suspects the marketing conversation has gotten too small for what is actually possible — and you are right.

The Signal Flow Map is the part most marketing teams have never drawn. The Meridian Readiness Checklist is the part most measurement consultancies will not tell you about, because the organizational checks are where their engagements quietly fail. The Cost Per Decision Worksheet is the part most AI vendors will never put on a slide, because it reframes their product from a productivity tool into a structural margin lever — which is a more honest sale, but a much harder one.

All three tools share a property: they connect things that are usually kept apart. The Google ecosystem and the proof layer. Marketing measurement and capital allocation. Decision velocity and the org chart. **That is the uncommon move.** Not better tactics. Better connections.

Where to go next

1. Pick the tool that maps to the conversation you are already losing — the executive briefing (Tool 1), the measurement budget request (Tool 2), or the AI investment justification (Tool 3) — and run it once this week.

2. Read the long-form thinking at uncommonmove.com/thinking for the deeper teardowns behind each tool — particularly the GA4 audit teardown, which is the prerequisite reading for the Meridian readiness checklist.

3. If you would like to talk about applying these to your specific organization, the best path is the newsletter — replies go to me directly. Find it at uncommonmove.com.

> *The uncommon move is the one most teams cannot afford to make — until the team that does makes it impossible to ignore.*

— Kyle Schwietz · Founder, Uncommon Move