The Unifying Message

How to sell one story to a buying committee that does not agree.

METHOD BY

David Levy

Principal, Broadfield LLC

PUBLISHED BY

Delivr.ai

29 August 2026

OPEN

The template is in step three.

No form, no download wall.

The Unifying Message

How to sell one story to a buying committee that does not agree.

METHOD BY
David Levy
Principal, Broadfield LLC

PUBLISHED BY
Delivr.ai
29 August 2026

The messaging matrix in step three is the working template. Take it.

In June 2025, Gartner published a number that should have stopped the personalization industry in its tracks. Across 1,464 buyers surveyed in North America, the UK, Australia and New Zealand, personalization produced a negative experience for 53% of them in their most recent purchase journey. Those buyers were 3.2 times more likely to regret the purchase, and 44% less likely to buy from the same brand again.

Read that again, because the usual defenses do not apply. This is not a finding that personalization is hard, or expensive, or slow to show returns. It is a finding that more than half the time, the thing we spent a decade building actively damaged the relationship it was supposed to deepen.

The industry response has been to buy more of it. Salesforce found that 84% of marketers admit to running generic, untargeted campaigns even while holding AI tools built to prevent exactly that, and that 98% of teams using AI report at least one data problem standing between them and personalization.

Something is broken that more software has not fixed. This piece argues it is not the software. It is that we have been personalizing to the wrong unit, and there is a specific, teachable method for personalizing to the right one.

SOURCES

Gartner, "Gartner Survey Reveals Personalization Can Triple the Likelihood of Customer Regret at Key Journey Points," 3 June 2025. Survey fielded November to December 2024 among a mixed population of B2B buyers and consumers; the regret figures are not broken out by segment.

Salesforce, *State of Marketing*, 10th edition, 2026, approximately 4,500 marketers across 26 countries.

Personalization survived the apocalypse it blamed

For most of the last five years, the field had an answer ready whenever results disappointed. We were losing our signals. The third-party cookie was going away, mobile identifiers were disappearing behind consent prompts, and precision was being regulated out of existence. Wait for the new plumbing, and relevance would return.

That story is now over, and not in the way anyone expected.

On 22 April 2025, Google announced it would not ship the standalone browser-wide prompt for third-party cookie choice, keeping its existing approach in Chrome. On 17 October 2025, it went considerably further and retired the Privacy Sandbox APIs themselves, including Topics, Protected Audience and Attribution Reporting, after six years of development. Third-party cookies remain in Chrome today with no removal date.

Meanwhile the mobile picture stabilized rather than collapsed. Adjust's benchmark puts global iOS ATT opt-in at roughly 38% in early 2026, up from about 35% a year earlier, following a flat trendline since 2023. Not a recovery, and not a collapse. A plateau.

So the apocalypse was called off. Marketers kept their cookies, kept roughly a third of mobile consent, and gained several years of additional first-party data tooling. If signal loss were the cause, results should have improved.

They did not. Which means the cause was somewhere else the whole time.

SOURCES

Google Privacy Sandbox, "Next steps for Privacy Sandbox and tracking protections in Chrome," 22 April 2025.
Search Engine Land and AdExchanger reporting on the retirement of the Privacy Sandbox APIs, 17 October 2025.
Adjust, "ATT opt-in rates: 2025 data and benchmarks," Mobile App Trends 2026, February 2026.

The gap nobody measures

There is a second number that gets less attention than it deserves. Deloitte Digital found that consumers recognize 43% of their experiences as personalized, while brands believe they are personalizing 61% of the time.

That 18-point gap is not a measurement error. It is the whole problem stated in miniature. Brands are counting the personalization they *emit*. Buyers are counting the personalization they *notice*. Those two numbers can drift apart indefinitely, because a message can be personalized on every dimension the sender tracks and still land as noise on the receiver.

The industry measures its own output and calls it performance. Nobody is measuring what the buying committee actually experienced, which is a different thing entirely and, in B2B, a much stranger one.

SOURCE

Deloitte Digital, *Personalization: It's a value exchange between brands and customers*, 2024.

You are not personalizing to a person. You are personalizing to a committee.

Here is the assumption buried under every personalization engine ever built: that the unit of personalization is an individual, and that relevance means matching a message to that individual's attributes and behavior.

In B2B that assumption is simply false, and it is false in a way that turns personalization from an asset into a liability.

Individuals do not buy B2B software. Committees do. A meaningful purchase runs through several people with different budgets, different risk exposure, different technical depth and different reasons to say no. They do not evaluate you in sequence. They evaluate you in a group chat you cannot see, comparing notes on what each of them was told.

Every personalization engine on the market optimizes each of those people independently, because each of them is a separate record. The engine is doing exactly what it was designed to do. The design assumed a buyer who decides alone.

Why more variants make it worse

Follow the standard playbook to its conclusion.

You have three personas on the committee. You are active in four channels. You segment by three journey stages. That is thirty-six message variants, and the better your tooling, the more confidently you will generate all of them.

Each variant is defensible on its own. Each was written to be maximally relevant to one reader at one moment. And each is therefore slightly wrong for the other two people on the committee, who will nonetheless see it, because committees forward things to each other.

The economic buyer hears about cost control. The technical evaluator hears about architecture. The champion hears about daily pain relief. All three are true. None of them is the same story. When those three people sit down together, they discover they have been sold three different products by a vendor who appears not to know its own value proposition.

This is what incoherence at the committee level looks like, and it is a much better explanation for the Gartner regret figure than any theory about weak content or bad data hygiene. The buyer is not overwhelmed because the messages were irrelevant. The buyer is overwhelmed because the messages were individually relevant and collectively incoherent.

Personalization, done the standard way, scales the incoherence.

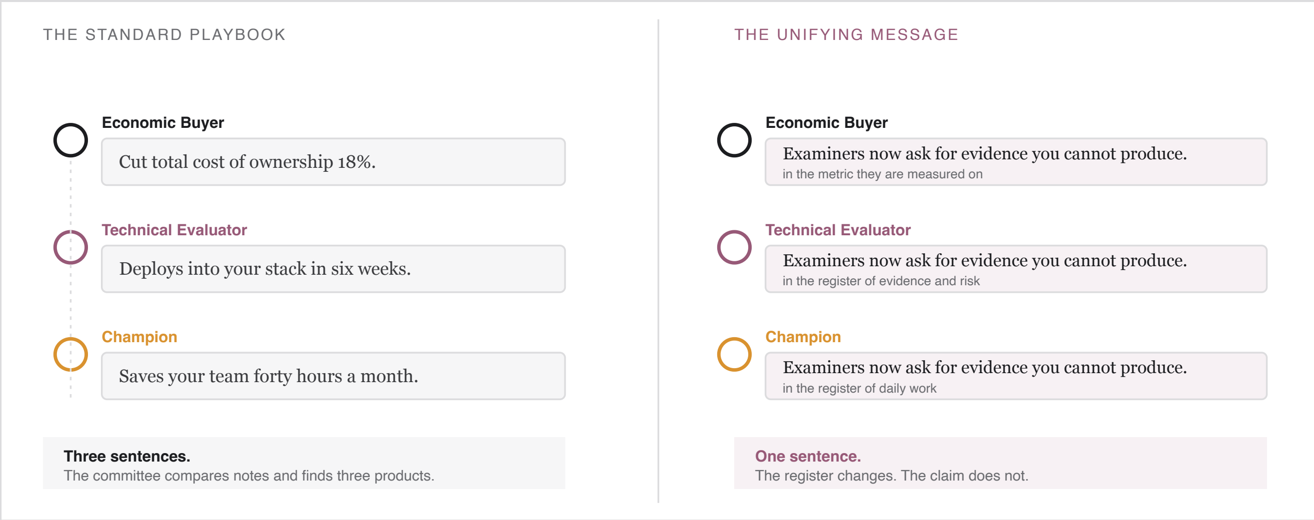

Both sides reach the same three people. Only the right-hand side survives those three people talking to each other.

The Unifying Message

The method described in the rest of this piece is David Levy’s. We run it inside every persona workshop we do, and it has been the most reliable predictor of whether a campaign survives contact with a real buying committee.

Its central move is the opposite of what the category teaches.

> A Unifying Message is the single sentence that must hold true for the economic buyer, the technical evaluator and the champion at the same time.

Not three sentences, tuned per persona. One sentence, constrained to survive three simultaneous readers who will compare notes. You find it before you write anything else, and everything else is downstream of it.

This is a constraint, not an expansion. The prevailing answer to a committee is to multiply messages until each member has one of their own. This method finds the invariant first, and treats variation as a matter of delivery rather than substance. Same message, different register, different channel, different evidence attached. One story.

The distinction that matters: unifying the *profile* is a data problem, and it has been solved several times over by good software. Unifying the *message* is a positioning problem, and no amount of profile unification produces it. A company can have a perfect single view of the customer and still tell three stories, because the three stories were written by three teams optimizing three funnels.

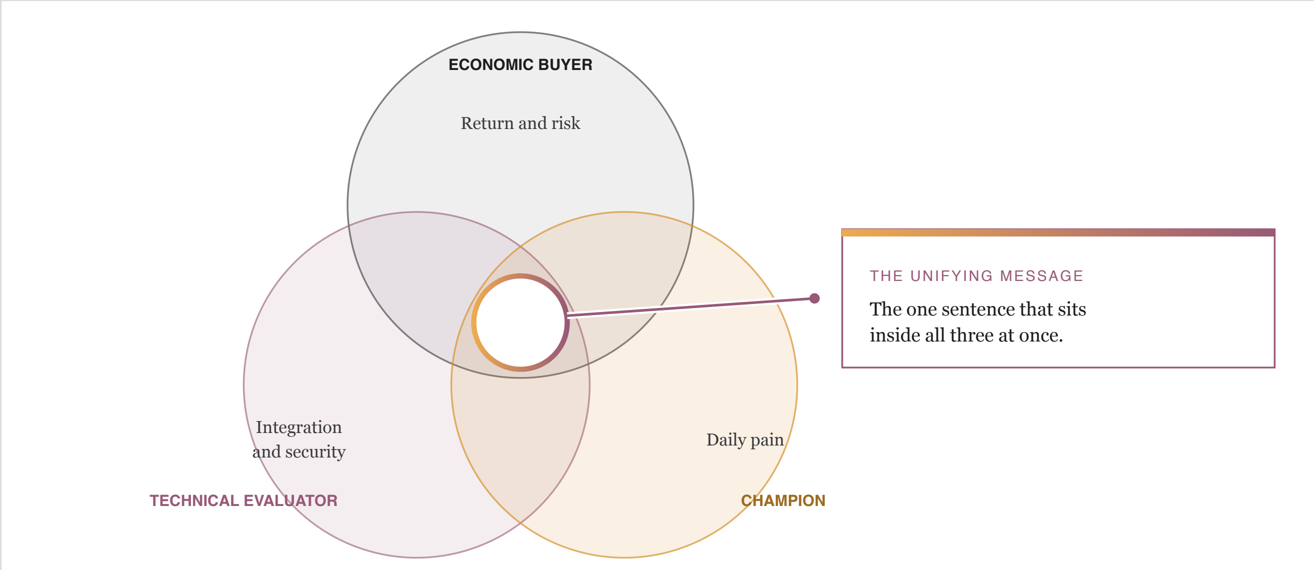

Not the average of three messages, and not three messages in a trench coat. The intersection, which is a smaller and much harder place to write from.

STEP 1 The persona workshop

Personas here are built from buying dynamics, not job titles.

A title tells you where someone sits. It does not tell you what they are accountable for when the quarter goes badly, what they were burned by last time, what they have to justify to their own boss, or what a failed purchase costs them personally. Those are the things that determine whether a message lands, and none of them is on an org chart.

The test of a persona built this way is whether it survives the matrix step later. Title-based personas collapse there, because when you try to write one sentence true for all three threads, title-based personas give you nothing to work with. They differ in seniority, not in stake. You end up writing a sentence true for their department rather than for them, which is the same as writing nothing.

STEP 2 Identify the buying group

Every deal runs three functional threads. Name a real title for each, per solution lane.

> **Economic Buyer, who signs**
>
> Controls budget, authorizes spend, owns the final go or no-go. Cares about return and risk, not features. If you cannot state your value in the metric this person is measured on, you do not have their attention regardless of how much they liked the demo.

> **Technical Evaluator, who builds the case**
>
> Vets integration, data fidelity, security and architecture. Can veto the entire deal on a single objection, and frequently does. This person is not looking for reasons to say yes. They are looking for the thing that will be their fault in eighteen months.

Champion, who surfaces the pain

Feels the problem daily, writes the tickets that eventually justify the budget line, and sells for you internally when you are not in the room. Most credible voice on the problem, and structurally unable to close it alone.

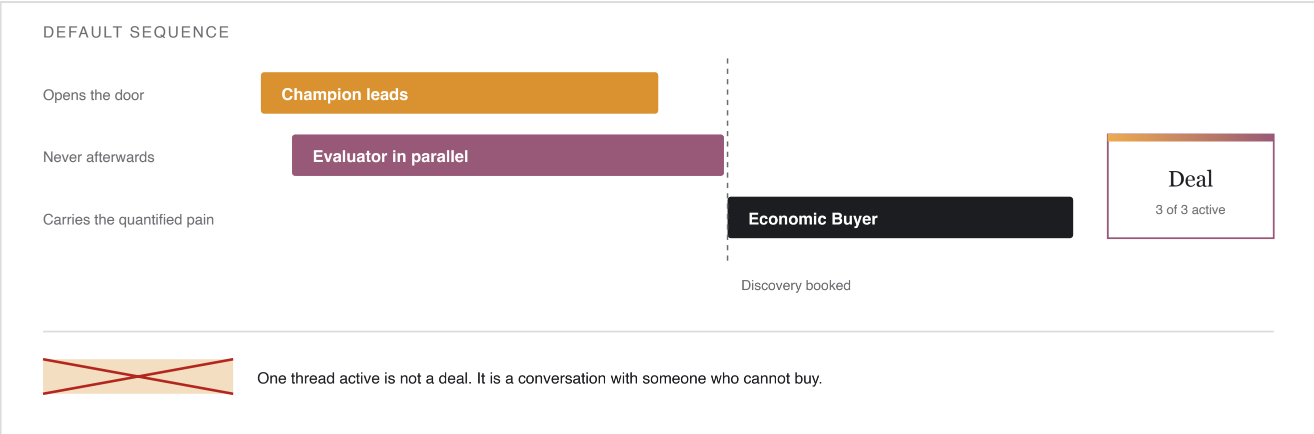

Order carries more of the outcome than most teams expect. A veto-holder who hears about the deal secondhand defaults to no.

When the committee is larger, layer the extended roles onto those three rather than inventing new threads. The Blocker or Skeptic was burned by this category before and distrusts it on principle, so neutralize early rather than late. The End User or Beneficiary rarely decides the purchase but determines the renewal. The Influencer shapes the decision without a vote, and may be an internal peer or an outside analyst.

Watch for lanes where a role collapses. In smaller organizations one person is routinely both Economic Buyer and Champion, which means selling to them once, in two registers, in the same conversation: the pain they feel and the number they answer for. Getting this wrong reads as either tone-deaf or condescending, and there is no recovery from either.

STEP 3 Build the messaging matrix

This is the deliverable, and it is one table.

One row per solution lane. A column for each thread, holding the named titles in that lane. And a final column, holding the Unifying Message. David's rule for that last cell is that it is the hardest and most important one in the entire document, and the matrix is not finished until it is right.

SOLUTION LANE	ECONOMIC BUYER (SIGNS)	TECHNICAL EVALUATOR (BUILDS THE CASE)	CHAMPION (SURFACES THE PAIN)	UNIFYING MESSAGE
Compliance automation, regional banks	Chief Risk Officer · VP Compliance	Director of Information Security · Security Architect	Compliance Analyst · BSA Officer	Examiners now ask for evidence your team cannot produce by hand, and the gap widens every quarter.

That example is a fictional composite, not a client lane.

Look at what the Unifying Message column is doing. It names a change in the world, not a capability in a product. It is true for the Chief Risk Officer, whose exposure grows as the gap grows. It is true for the Security Architect, who is the one who will have to produce that evidence from systems that were never designed to. And it is true for the Compliance Analyst, who is already doing it by hand and knows exactly how far behind they are.

Most first attempts fail, and they fail in a specific and predictable way. They come out true for the Champion and merely tolerable for everyone else. Here is a typical failed candidate for the same lane:

> *Automate your compliance workflows and save your team forty hours a month.*

The Champion loves it, because it is their forty hours. The Chief Risk Officer is indifferent, because forty analyst hours is not a number that appears anywhere in their risk register. The Security Architect is actively suspicious, because "automate" is the word that precedes every integration that later becomes their problem. One sentence, one enthusiast, one shrug and one raised eyebrow. That deal is now single-threaded and nobody has noticed yet.

Three tests for a candidate sentence:

1. **Would each thread repeat it out loud to their own boss, in their own words?** If the Economic Buyer would not say it in a board meeting, it fails, no matter how well it tests with the Champion.
2. **Does it name a change in the world rather than a feature of your product?** Features are true about you. Changes are true about them. Only the second kind survives three readers.
3. **Would any thread be surprised to learn the others received the same sentence?** If yes, you have not written a Unifying Message. You have written three messages and a coincidence.

STEP 4 Sequencing and the multi-thread rule

The matrix tells you what to say. It does not tell you what order to say it in, and order carries more of the outcome than most teams expect.

The sequence David teaches: the Champion leads, because they feel the pain daily and will open the door. The Technical Evaluator runs in parallel, not afterwards. This is the part teams get wrong most often. A veto-holder who first hears about your deal secondhand, from a colleague who is already enthusiastic, defaults to no. Warm them before they are surprised. The Economic Buyer engages after discovery is booked, and engages armed with the Champion's pain already quantified in the Economic Buyer's own metric.

Triggers invert this. A public incident makes the Economic Buyer the one already searching, so go executive-first and do not make them wait for a bottom-up motion. A dated compliance deadline makes the Technical Evaluator the initiator, because they own the deadline.

And the rule that governs all of it: any deal with only one active thread is single-threaded risk, whatever the CRM says about the stage. One thread is not a deal. It is a conversation with someone who cannot buy.

The method has a prerequisite nobody talks about

Everything above is executable on a whiteboard. That is what makes it teachable, and it is why the workshop works.

But look closely at what step four actually requires you to do, and a prerequisite appears that no amount of workshop facilitation supplies.

The matrix says three named threads at one account must each receive the same message in their own register, in a coordinated order, within the same window of time. To do that, you have to know which anonymous session on your site belongs to the Evaluator and which belongs to the Champion, and you have to know they work at the same company, and you have to know it this week rather than next quarter.

Absent that, the matrix is an artifact. A good one, worth the workshop, and it will make your messaging measurably more coherent wherever a human is manually driving the account. But it will not run as a campaign, because campaigns execute against records, and the records are the problem. Your CRM knows the Champion, who filled in a form. It does not know the Evaluator, who read four technical pages last Tuesday and told nobody.

This is where identity resolution stops being an infrastructure concern and becomes a messaging concern. Resolving anonymous traffic to known people and known accounts is what converts the matrix from a document into a targetable plan. At Delivr.ai, that resolution is 100% deterministic, meaning matches come from direct identifier linkage rather than modeled inference. For a method built on saying the same true thing to three specific people, inference is not good enough. If the Evaluator is a probability, the sequence in step four is a guess.

The compliant posture is the same posture

There is a version of this argument that gets a company sued, and it is worth separating it from the version that does not.

The legal environment around web tracking has tightened sharply. Twenty states now have comprehensive consumer privacy laws in effect. California's AB 566, the Opt Me Out Act, signed 8 October 2025, will require browsers themselves to ship a native opt-out signal from 1 January 2027, escalating an existing duty to honor such signals into a requirement that the mechanism exist by default. Global Privacy Control is treated as a binding opt-out by regulators in California, Colorado and Connecticut, who ran a joint compliance sweep in 2025.

Litigation has followed. An industry litigation tracker reports growth in California wiretapping-style claims over website trackers from 54 filings in 2022 to several thousand pending cases by mid-2026. Legislative testimony in support of California's SB 690 cited growth from roughly 600 filings to more than 4,000 under the relevant provision. These are commercial and legislative counts rather than official court statistics, and should be read as directional.

The point for this method is narrow and practical. A Unifying Message strategy needs to know who three specific people are. It does not need to know everything about them, and it does not need signal obtained in ways that create exposure. Consented, deterministic, person-level resolution is both the more accurate route and the more defensible one. Inferred and scraped signal is where both the accuracy problem and the litigation surface live.

For the statutory analysis in depth, see our separate piece on the legality of person-based intent data.

SOURCES

MultiState, "20 State Privacy Laws in Effect in 2026," 4 February 2026.

IAPP, "California governor signs new law requiring in-browser opt-out preference signal," 9 October 2025.

Cookie-Script CIPA Lawsuit Tracker. Bloomberg Law reporting on July 2026 California Assembly Privacy and Consumer Protection Committee testimony.

Why this page is not gated, and neither is the template

The matrix template above is the actual working artifact. It is not a teaser for a gated version with the good parts removed.

We are an identity resolution company. Putting a form in front of this page, to collect an email address from a reader we are able to resolve anyway, would be an admission that our own product does not work. Gating is what vendors do when they cannot identify their own traffic.

Take the template. Run the workshop yourself. If the method works and you eventually need to execute it against people who never filled in a form, you will know where to find us.

Run it Monday

1. **Pick your highest-value solution lane and name the three threads.** Real titles, not categories. If you cannot name a Technical Evaluator, that is not a lane without one. That is a veto-holder you have not met.
2. **Write the Unifying Message and run the three tests.** Expect the first attempt to fail test one. Most do. Rewrite until a sentence survives all three readers, then check it against your last four campaigns and count how many of them contradicted it.
3. **Find out how many of the three threads you can actually reach.** Not how many are in your CRM. How many you could put a coordinated sequence in front of this week, at the same account. That number is the real ceiling on the method, and for most teams it is considerably lower than they assume.

The first two steps cost you an afternoon. The third one is the one that tells you whether you have a messaging problem or an identity problem. Most teams discover they have been trying to fix the first while the second was quietly capping every campaign they ran.

METHOD BY

David Levy

PRINCIPAL, BROADFIELD LLC · MARKETING AND STRATEGY ADVISORY

David Levy is Principal of Broadfield LLC, a marketing and strategy advisory practice working with owners, CEOs and marketing leaders on positioning, go-to-market and marketing operating models. He spent nearly six years at Foundever, most recently as SVP, Global Revenue Marketing, where he stood up the global revenue marketing function and built a 1:1 account-based tier for top US and EMEA accounts that moved shared sales and marketing targets away from lead volume and toward account-level coverage. Earlier he served as the company's US CMO and, before that, built its US marketing organization from near zero following acquisition.

His agency career includes global strategy leadership at MRM on Cisco's demand generation and digital innovation account, where persona-driven content personalization and engagement scoring were central, along with earlier roles at VML, DEG and Torch.AI.

The Unifying Message comes out of that work. Twenty-five years of enterprise go-to-market, and one finding that kept repeating: deals turn on whether the buying committee heard one story or three.

LINKEDIN PROFILE →

The Unifying Message method, the three-thread buying group model and the messaging matrix are David Levy's work. Delivr.ai runs the method inside its persona workshops, and this piece is its first published description. The three tests in step three are our own addition to it.