

The Intelligent ICP

From Static Profiles to Dynamic Opportunity

MOST ICPs ARE STATIC. YOUR BEST CUSTOMERS AREN'T.

A practical guide to finding, understanding, and engaging your best customers in a more intelligent way.



See your market more clearly.



CONNECT
More Signals



UNDERSTAND
Deeper Context



DECIDE
With Confidence



ACT
At the Right Time



MEASURE
What Matters

01 A CLEARER PATH FOR WHAT'S NEXT

Executive Summary

Most ICPs are built to describe the past.
Intelligent ICPs help you act in the present.

A new approach to the ideal customer profile helps you see more signal, move faster, and create more value.

CLEARER
INSIGHTS
BRIGHTER
OPPORTUNITY



01

The problem

Most companies still define their ideal customer profile using static characteristics like industry, company size, revenue, geography, or title. That approach can be useful, but it does not tell a business who deserves attention now.



02

What changed

Businesses now have access to far more signal than they used to. Behavior, purchase history, engagement, timing, intent, service activity, and external context can all improve how a company identifies the right opportunity.



03

What this guide argues

An ICP should no longer function as a frozen marketing persona. It should function as a living decision system that helps teams recognize fit, understand context, prioritize effort, and decide the next best action.



04

What executives should do next

Start with one practical question: Which customers are most likely to create value now? Then connect the signals you already have, define the decision you want to improve, run a 30-day test, and measure the business outcome.

IN ONE SENTENCE

The old way described your target.
The intelligent way helps your business identify, prioritize, and act.



CONNECT

Gather better signals



UNDERSTAND

See context



DECIDE

Prioritize the right opportunity



ACT

Move with precision



MEASURE

Track the outcome

READ THIS FIRST

A clearer path. A practical guide.

Use the guide to improve one decision: which account deserves attention, for what reason, and with what next step?

UNDERSTAND	The old way, the new model and what “best” means.	04–07
BUILD	Gather evidence. Separate fit from timing. Create an explainable queue.	08–15
PROVE	Apply it, price the effort, test fairly and protect the customer.	16–22
USE	Complete the worksheets. Check the sources. Make the next decision.	23–26

HOW TO READ THE EVIDENCE

Established guidance: modern ICP practice already considers customer behavior and market context [1]. The contrast here is a static workflow versus a connected one, not “old marketing was wrong.”

TAG framework: the matrices, weights, thresholds and 30-day plan are recommended starting points. Examples are fictional; numbers are illustrative. They are not client results or performance promises.

01 / THE SHIFT

Keep the fundamentals. Upgrade the workflow.

A profile is useful. A profile connected to evidence, ownership and action is useful at the moment of a decision.

BUSINESS JOB	STATIC IMPLEMENTATION	INTELLIGENT IMPLEMENTATION
Define	Industry, size and geography on a slide.	Keep those filters. Add the problem solved, economics and ability to deliver.
Research	A periodic workshop and a purchased list.	Combine customer outcomes, interviews and permitted external evidence.
Prioritize	Largest account or most recent click wins.	Separate structural fit, current need and evidence quality.
Act	Hand over a list and hope it gets worked.	Attach a reason, next action, owner and review date.
Learn	Count leads and celebrate activity.	Track completed actions, qualified progress and contribution.
Refresh	Revisit the profile when the plan changes.	Refresh triggers as needed; revise the profile when outcomes justify it.

KEEP THIS

Customer interviews, industry expertise and firmographics still matter. Intelligence adds discipline; it does not replace them.

02 / THE DEFINITIONS

Fit is not timing. A person is not an account.

Use different tools for different questions. Mixing them into one mysterious score hides the decision you actually need to make.

ICP

Which companies fit?

The type of business likely to gain value from your offer, work well with your delivery model and create a worthwhile relationship.

BUYER PERSONA

Who is involved?

The people inside the account: the user, champion, budget owner, approver and others who influence the purchase.

ACCOUNT SIGNAL

What changed?

A dated observation: a quote request, renewal window, service issue or confirmed expansion. It is evidence to interpret, not proof of intent.

NEXT BEST ACTION

What should happen?

The recommended move for that account now: qualify, schedule, nurture, investigate, resolve an issue or decline.

THE DISTINCTION

The ICP defines where you should compete. Signals help decide where to spend attention today.

Source basis: ICP and persona definitions [1]; separate fit and engagement scoring [2].

03 / DEFINE "BEST"

Your biggest customer may not be your best.

Begin with value delivered to the customer and value retained by your business. Revenue alone is an incomplete target.

FICTIONAL 12-MONTH CUSTOMER COMPARISON

Revenue

\$ thousands

Account A \$120k



Account B \$80k



Contribution

\$ thousands

Account A \$12k



Account B \$24k



SAME 12-MONTH WINDOW	ACCOUNT A	ACCOUNT B
Revenue	\$120,000	\$80,000
Gross profit	\$30,000	\$32,000
Additional cost to serve	(\$18,000)	(\$8,000)
Contribution	\$12,000	\$24,000

Define the costs once. Contribution here means gross profit less additional customer-specific servicing costs not already included in gross profit. It excludes corporate overhead and tax. Also inspect retention, payment behavior and whether the customer achieved the promised outcome.

04 / BUILD THE PROFILE

Learn from the customers you won. And lost.

A useful profile needs a comparison group, a common time window and a reason the customer chose you.

WON + HEALTHY

Value achieved.
Worth serving.

WON + STRAINED

High effort.
Thin contribution.

LOST + NO DECISION

Why did buying
not happen?

EMERGING

Different on paper.
Promising in practice.

1 Choose a comparable window.

Compare accounts with enough time to buy, onboard and show an outcome. Do not compare a new account's first month with a mature account's full year.

2 Join the minimum useful records.

Match CRM account IDs to transactions, delivery effort and conversations. Keep unknown costs and missing records visible rather than filling them with guesses.

3 Ask what changed the decision.

Interview customers and lost prospects: What triggered the search? What alternative did you consider? Why act then? What nearly stopped you? What value appeared after purchase?

START SMALL

A few varied interviews can challenge assumptions. They are a discovery sample, not proof that the pattern represents the whole market.

Source basis: begin with customer outcomes, interviews and team feedback [1]. Cohort and interview method: TAG recommendation.

05 / THE EVIDENCE LADDER

More signals. Not more assumptions.

Give each observation a source, date and account match. Keep what you observed separate from what you inferred.

SIGNAL CLASS	EXAMPLES	APPROPRIATE USE
Direct customer evidence	Quote request, stated deadline, agreed next step.	Confirm scope and readiness with the customer.
Internal operating evidence	Purchase history, repeated service issue, expiring agreement.	Check context, relationship health and data completeness.
Primary external evidence	Company announcement, official tender, published expansion.	Verify the entity and event date; investigate relevance.
Indirect evidence	Aggregated intent, third-party enrichment, anonymous activity.	Use as a research lead. Do not treat it as a buying decision.

MINIMUM EVIDENCE RECORD

Account ID + observed fact + source link or record ID + event date + capture date + confidence + expiry/review date + permitted use + owner.


Unknown is a valid state. Missing evidence does not mean poor fit, no need or permission to invent an answer.

TAG evidence hierarchy and operating rules. Reliability, privacy and accountability principles: NIST [5].

06 / INTELLIGENCE GATHERING

Gather. Verify. Then interpret.

Search around a business question, not around a desire to collect everything. A new location is a fact; a need for your service is a hypothesis.



ILLUSTRATIVE RESEARCH NOTE

Northstar Field Services

Observed: an official company announcement describes a second location. The source date and account identity are recorded.

Inferred: coordinating inquiries across locations may be harder. There is no verified evidence of missed calls or budget.

Next step: ask the operations contact how inquiries are routed today. Do not claim to know they have a problem.

SOURCE RULE

Two articles repeating the same announcement are one underlying source, not two independent confirmations.

07 / FRESHNESS

Two clocks. One clearer decision.

Customer fit usually changes for different reasons than buying readiness. Refresh each on the cadence its decision requires.

THE SLOWER CLOCK


Structural fit

Problem, offer, operating model and economics. Review after meaningful changes and in a scheduled business review.

THE FASTER CLOCK

Current readiness

Requests, deadlines and active conversations. Refresh when events arrive; check whether the need is still open.

ILLUSTRATIVE EVENT WEIGHT, NOT A CONVERSION PROBABILITY**DO NOT CONFUSE**


An old click can lose relevance. A service restriction, active opt-out or unresolved issue does not expire because a score decays.

Separate fit and engagement, including event decay, are supported in HubSpot [2]. Cadences below are TAG starting points, not universal rules.

08 / THE PRIORITY MATRIX

Right fit. Right time. Different actions.

Use the matrix only after checking permission, serviceability and evidence quality. Urgent inbound requests still receive normal service.



Missing or contradictory evidence? Send the account to research, not automatically to the bottom of the list.

Capacity matters. Match the daily queue to what the team can actually work. Ten owned actions are more useful than a hundred alerts with no follow-through.

09 / A TRANSPARENT STARTER MODEL

Make the score explain itself.

Start with a small, inspectable model. Keep fit, readiness and confidence separate. A score of 90 is not a 90% chance of winning.

DIMENSION	QUESTION TO EVIDENCE	WEIGHT
Problem fit	Does this account have a problem the offer is designed to solve?	30%
Delivery fit	Can you serve its geography, scale, systems and requirements?	25%
Economic fit	Is the expected relationship worth acquiring and supporting?	25%
Success fit	Can the customer adopt the solution and realize the outcome?	20%

Fit score = sum of (dimension rating ÷ 5 × weight).

Rate each dimension from 0 to 5 with a written reason. Zero means verified poor fit; 3 means partial fit; 5 means strong fit. Unknown stays unknown.

Starter routing: 75–100 = strong fit; 50–74 = review; below 50 = weak fit. These are test settings, not industry benchmarks. Apply hard exclusions before any score. Never turn an opt-out into a small negative point adjustment.

Predictive is a later option. Microsoft's product, for example, requires at least 40 qualified and 40 disqualified leads for lead scoring. That is a product minimum, not proof of sufficient data for your use case [3].

TAG weights and thresholds are illustrative. Rule-based scoring concepts [2]; labeled outcome requirements for predictive tools [3].

10 / WORKED EXAMPLE

Three accounts. Three different moves.

Fictional accounts for a B2B service provider. The same fit formula produces different priorities once timing and evidence are visible.

ACCOUNT	RATINGS ¹	FIT	READINESS EVIDENCE	NEXT ACTION
Northstar Field Services	5 / 4 / 4 / 5	90	Requested a relevant discovery call yesterday.	Prioritize. Confirm the problem and book the next step.
Summit Service Group	5 / 5 / 4 / 4	91	Strong fit; no verified near-term need.	Nurture. Share relevant guidance and set a review date.
Atlas Enterprise	1 / 1 / 2 / 3	33	Active interest, but needs exceed delivery scope.	Qualify or refer. Do not overpromise to win a meeting.

¹ Rating order: problem, delivery, economics, success. All example ratings are assumed known. Evidence confidence is high only after the source, date and account identity have been checked.

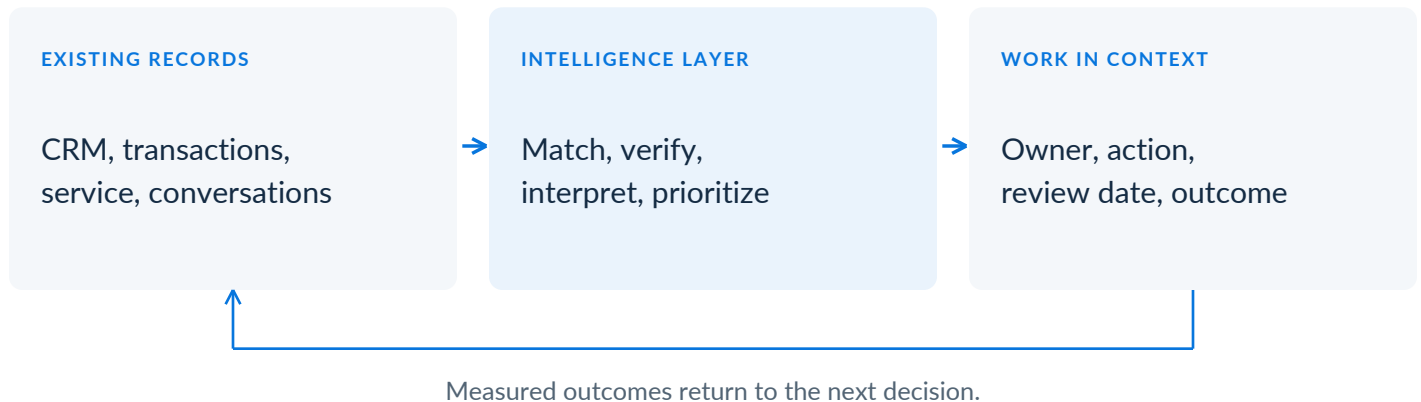
WHY NORTHSTAR GOES FIRST

It is not the highest fit score. It combines sufficient fit with a direct, current request. The decision is explainable without pretending to predict the future.

11 / THE OPERATING SYSTEM

Connect what you have. Add only what helps.

The goal is an evidence-backed action queue inside the workflow people already use, not another destination they must remember to check.



USE THE SIMPLEST TOOL	PUT IT TO WORK
Rules and code	Account matching, arithmetic, exclusion checks, permissions, duplicate prevention and approved routing.
Models where useful	Extract a stated need, summarize a call, classify a problem or propose a next question. Retain supporting evidence.
People where necessary	Resolve conflicting facts, approve sensitive outreach, change policy and make commitments to customers.

Minimum action record: account ID, fit dimensions, readiness, confidence, sources, policy checks, proposed action, owner, due date, outcome and rule/model version.

Architecture and control design: TAG recommendation. Simple workflows before additional agentic complexity: Anthropic [4].

12 / THE DAILY WORK QUEUE

A priority is useful only when someone owns it.

Deliver a short action brief, not a dashboard full of scores. Make the next step easy to understand, accept, change or reject.

ILLUSTRATIVE ACCOUNT BRIEF

Northstar Field Services

FIT

90 / 100

READINESS

Direct request

CONFIDENCE

Verified

Why now: requested a discovery call yesterday about inquiry handling across two locations. The contact and request match the account.

What to do: confirm the current handoff process and arrange a 20-minute discovery conversation. Do not assume a missed-revenue amount.

Who / when: assigned account owner; respond by the team's agreed inbound service deadline. **Evidence:** inquiry record + dated call note. **Check:** contact permission and scope.

ACCEPT

REVISE

RESEARCH

REJECT + REASON

Close the loop. Record the result: scheduled, not ready, wrong contact, no fit, blocked or completed. A manager reviews patterns in rejected recommendations, not just totals.

13 / PRACTICAL APPLICATIONS

One model. Several business moments.

Use separate criteria for acquisition, expansion and retention. A prospect's buying signal is not the same as a customer's service risk.

MOMENT	RELEVANT EVIDENCE	NEXT ACTION	OUTCOME TO WATCH
Acquire B2B services	A verified operating change plus a problem you can serve.	Ask a relevant discovery question.	Qualified next steps per eligible account.
Expand Distribution	A customer buys one category; a relevant adjacent product is available.	Check need, stock and contribution before recommending.	Incremental contribution, not just order value.
Retain Recurring services	An unresolved issue or an approaching renewal changes the relationship.	Resolve the issue before adding a sales pitch.	Issue resolution and renewal outcomes.
Respond Local service teams	A direct inquiry states a service need, location and timing.	Confirm serviceability; route or schedule appropriately.	Completed qualified appointments and response time.

BOUNDARIES

These are workflow examples, not sector benchmarks. Adapt to delivery capacity, permitted data and the obligations of the business.

14 / A SIMPLE BUSINESS CASE

Make the economics visible before you scale.

Use one evaluation window and count only incremental value. The table below is an illustrative sensitivity test, not a forecast.

90-DAY ASSUMPTION

\$2,000 per extra win

Contribution realized within the window, after delivery and customer-specific costs.

90-DAY ASSUMPTION

\$3,000 program cost

Setup, tools, data, review, outreach effort and monitoring for the same test window.

INCREMENTAL WINS	ADDED CONTRIBUTION	NET AFTER PROGRAM COST
0	\$0	-\$3,000
1	\$2,000	-\$1,000
2	\$4,000	+\$1,000
4	\$8,000	+\$5,000

Break-even: $\$3,000 \div \$2,000 = 1.5$, so at least **2 incremental wins** under these assumptions. Actual value and cost may be different.

Do not double-count. Saved hours are capacity, not cash savings unless the capacity is productively redeployed or expense is actually removed. Revenue, contribution and cash receipts are different measures. Count wins over a credible comparison, not every sale touched by the system.

15 / PROOF DESIGN


Test the new method. Not your optimism.

A high score can correlate with buying without proving your workflow caused improvement. Separate model evaluation from a live business test.

OFFLINE CHECK

Replay past decisions using only evidence that existed at that moment. Keep later outcomes out of the inputs, and evaluate on a later, untouched period. Otherwise the model can appear accurate by seeing the answer in advance [6].

LIVE COMPARISON



Compare all assigned accounts, not only contacted accounts. If each pool has 100 accounts and capacity for 50 contacts, keep 100 in each denominator. This tests prioritization under equal capacity. The counts are an example, not a sample-size recommendation.

Report uncertainty. Agree the outcome, duration and analysis before the test. Low volume or a long sales cycle may require more time. Without a credible control, describe an operational improvement, not proven incremental revenue [7].

Source basis: data leakage [6]; randomization, statistical power and evaluation limits [7]. Pilot adaptation: TAG recommendation.

16 / THE FIRST PROOF

Thirty days to a better-informed decision.

A 30-day sprint should establish usability, reliability and early progress. It may not be long enough to establish revenue impact.

WHEN	WORK	DELIVERABLE	OWNER
Days 1–7 Define	Choose one offer, segment and business decision. Agree the baseline, exclusions, outcome and owner.	A one-page test charter.	Business sponsor
Days 8–14 Prepare	Match records, review example accounts, check evidence and tune simple rules in shadow mode.	A reviewed sample and an explainable queue.	Data owner + sales lead
Days 15–21 Use	Run the approved comparison. Give both groups equal capacity; log actions, overrides and defects.	Real use with an auditable action trail.	Account owners
Days 22–30 Review	Compare early outcomes and cost. Fix issues. Decide whether to stop, revise or continue observation.	A decision memo, not a victory slide.	Sponsor + analyst

AFTER DAY 30

Follow longer-cycle opportunities to agreed day-60 or day-90 checkpoints. Do not label meetings or pipeline value as realized profit.

17 / EXECUTIVE SCORECARD

Measure the result. Keep the guardrails.

Pick one primary business outcome, a few diagnostic measures and explicit stop conditions. Define every denominator before launch.

LENS	MEASURE	INTERPRETATION
Business progress	Accounts reaching an agreed qualified next step ÷ all eligible assigned accounts.	Leading indicator; does not equal revenue.
Commercial result	Observed contribution per assigned account and per selling hour, over the same window.	Separate realized value from estimates.
Execution	Completed assigned actions ÷ actions due; time from trigger to appropriate response.	A good queue still needs follow-through.
Recommendation quality	Unsupported or wrong priorities ÷ independently audited recommendations.	Audit beyond the highest-scoring accounts.
Customer protection	Suppression failures, unapproved actions and serious data-handling incidents.	A serious breach triggers pause and review.

SCALE	REVISE	STOP
Evidence is sufficient; value and controls hold.	Useful direction; fixable gaps remain.	Unsafe, unusable or not economically justified.

TAG operating scorecard. General experimental measurement and interpretation principles: [7].

18 / TRUST AND HUMAN CONTROL

Intelligent. Accountable by design.

A recommendation should be traceable, contestable and safe to ignore. No score earns permission to bypass policy.

CONTROL	PRACTICAL RULE	OWNER
Permitted data	Use approved sources and purposes. Minimize personal data. Set access and retention rules; exclude sensitive personal attributes from routine commercial prioritization.	Data owner
Evidence before action	Retain source, date, identity match and uncertainty. Unsupported model statements go to review, not into the customer record as fact.	Account owner
Hard stops	Enforce opt-outs, channel restrictions, service limits and required approvals independently of the score. A high score cannot cancel a prohibition.	Workflow owner
Human judgment	Require approval for consequential commitments, ambiguous recommendations and changes to targeting policy. Let users override with a reason.	Sales leader
Bias and feedback	Review neglected segments, false negatives and uneven data coverage. Historical buying patterns may reflect who was previously contacted, not who could benefit.	Sponsor + analyst

Scope: this guide addresses B2B commercial prioritization, not automated eligibility, underwriting, hiring or other high-impact decisions. Check applicable obligations and sector requirements before deploying.

TAG safeguards informed by NIST trustworthiness guidance [5] and the voluntary AI RMF Playbook [8]. Not legal advice.

19 / THE TAG JOURNEY

Start where you are. Expand what proves useful.

Do not require a company-wide transformation before improving one decision. Add the next data domain when it changes an important action.

START**One decision.**

Use a manageable account set, existing records and a simple fit model.
Name the owner and baseline.

PROVE**One useful workflow.**

Deliver evidence-backed actions. Measure use, quality, economics and customer impact.

EXPAND**One meaningful addition.**

Add cost, stock, payment or service data only when it improves the next decision.

AT EVERY STAGE

CONNECT → **UNDERSTAND** → **DECIDE** → **ACT** → **MEASURE**

Act and Measure begin with the first module. Learning is part of delivery, not a final phase.

WORKSHEET 01 / FILL IN THE PDF

Your intelligent ICP canvas.

Complete this with the people who sell, deliver and support the offer. Record facts separately from hypotheses.

OFFER / BUSINESS SEGMENT

OWNER / REVIEW DATE

WHAT MEANINGFUL OUTCOME DOES THE CUSTOMER NEED?

FIT CRITERIA + EVIDENCE THAT SUPPORTS THEM

HARD EXCLUSIONS / UNKNOWN INFORMATION

WHICH CURRENT SIGNALS JUSTIFY ATTENTION?

NEXT ACTION / OWNER / DUE DATE

SUCCESS MEASURE / BASELINE / REVIEW WINDOW

TIP: "A large company" is a description. "A serviceable account with a verified problem and viable economics" is a decision rule.

WORKSHEET 02 / FILL IN THE PDF

Your first proof charter.

Agree what success, uncertainty and failure will look like before the first live recommendation reaches the team.

DECISION TO IMPROVE / HYPOTHESIS

ELIGIBLE ACCOUNT POOL / EXCLUSIONS

BASELINE / COMPARISON METHOD

PRIMARY OUTCOME / DENOMINATOR

ALL-IN COST / VALUE WINDOW

GUARDRAILS / APPROVALS / STOP CONDITIONS

SPONSOR / OPERATOR / DATA OWNER

START / DAY-30 / LATER CHECKPOINT

RESULT / UNCERTAINTY / SCALE, REVISE OR STOP

Keep the signed-off version. Log changes during the test. A changed method may need a new comparison rather than a new success story.

RESEARCH AND EDITORIAL NOTES

Sources. Method. Limits.

Sources checked September 24, 2026. Linked titles open the original material. Vendor documentation establishes capability, not guaranteed business impact.

- 01** **Salesforce** · May 14, 2026
[Ideal customer profiles: definition and development](#)
Definitions, customer evidence and iterative refinement. Used on pp. 3, 5 and 7.
- 02** **HubSpot** · Updated September 2, 2026
[Overview of the lead scoring tool](#)
Separate fit and engagement scores, event limits and score decay. Used on pp. 5, 10 and 12.
- 03** **Microsoft Learn** · Updated February 27, 2026
[Lead and opportunity scoring](#)
Predictive scoring depends on labeled historical outcomes. Used on p. 12.
- 04** **Anthropic** · December 19, 2024
[Building effective agents](#)
Use simple workflows before adding agentic complexity. Used on p. 14.
- 05** **NIST** · Living resource
[AI Risk Management Framework FAQs](#)
Trustworthiness includes reliability, privacy and accountability. Used on pp. 8 and 21.
- 06** **scikit-learn** · Living documentation
[Common pitfalls and recommended practices](#)
Prevent leakage from future information and test data. Used on p. 18.
- 07** **Kohavi, Longbotham, Sommerfield & Henne** · 2009; online July 30, 2008
[Controlled experiments on the web: survey and practical guide](#)
Randomization, sample size and trustworthy evaluation. Used on pp. 18 and 20.
- 08** **NIST** · AI RMF 1.0 companion
[AI RMF Playbook](#)
Voluntary actions for Govern, Map, Measure and Manage. Used on p. 21.

Method and limits. The cover and opening summary are editorial framing, not survey estimates of static-ICP adoption. "Intelligent ICP" names TAG's recommended workflow. Fictional examples, weights and economic scenarios are not empirical results; test them locally. No independent validation of this complete framework is claimed.

YOUR NEXT MOVE

Know who matters. Act when it matters.

The value is not a more sophisticated profile. It is a better decision, acted on responsibly, with an outcome you can measure.

01

Choose the decision.

One offer. One account segment. One outcome that matters.

02

Connect the evidence.

Use what you have. Label what you know. Expose what is missing.

03

Prove the action.

An owner, a next step, a fair comparison and a review date.

THE WORKSHOP QUESTION

Which ten accounts deserve a different action this week, and what evidence would justify it?