Framework 97 min read April 2026 Updated 3 October 2026 By Charlton Jones

The TRACE Framework: From Click to Funded Client

Truth · Route · Activate · Churn Engineer · Economics. A practitioner's framework built from walking the funnels of more than 30 retail FX and CFD brokerages end to end and measuring exactly where acquisition spend disappears before it becomes funded-account value. Published here in full, as a build specification.

Jupiter Tech Journal · https://jupitertech.io/blog/trace-framework · Markdown version: https://jupitertech.io/blog/trace-framework.md

Regulatory notice. TRACE is an operational framework. It is not legal or compliance advice. Before deploying any component involving tracking, profiling, automated decisions or outbound messaging, obtain compliance and legal sign-off on lawful basis, consent, cross-border data transfers and applicable financial promotion rules. Appendix E is the compliance readiness checklist.

Key takeaways

- TRACE is a five-stage revenue engineering framework for retail FX and CFD brokerages: Truth, Route, Activate, Churn Engineer, Economics. It is a deployment order, not a tool, and each stage enables the next.
- Most brokerages do not have a lead problem; they have a leakage and allocation problem, and buying more traffic makes it worse.
- Three fixable leakage points, speed to lead, KYC drop-off and failed deposits, cost the benchmark 1,000-lead-per-month brokerage roughly $600,000 to $900,000 a year. The model and every calculation are shown in full.
- Only three of the 16 components need machine learning. The 13 deterministic automations ship first, as rules, and deliver most of the measurable ROI.
- No CRM migration is required. Ownership of the event stream is.
- The rollout is 12 build weeks after a mandatory Week 0 of prerequisites, about 90 days, and a small team that completes Phase 1 fully recovers more revenue than a larger team that half-ships everything.

BASED ON AUDITS OF 30+ FX/CFD BROKERAGES

What TRACE is, who it is for, and the benchmark used

TRACE is a five-stage revenue engineering framework for retail FX and CFD brokerages: **Truth, Route, Activate, Churn Engineer, Economics**. It sequences 16 event-triggered components, three AI agents and 13 deterministic automations, so that a brokerage recovers the acquisition spend that leaks between an ad click and a funded, trading client. Jupiter Tech built it from direct operational audits of more than 30 brokerages, each run as a prospect: registering, completing KYC, attempting deposits, trading on MT4 and MT5, requesting withdrawals, and tracing the events behind every step to see where spend disappears before it becomes funded-account value. The audits were run between 2024 and 2026 across more than 30 retail FX and CFD brokerages of different sizes and licences; the benchmark figures sit at values that recur across that set.

The framework tells you which components need machine learning and which do not, and why the automations deliver most of the measurable ROI. It is designed to be read by a Head of Sales, a COO or a CTO and used as a build specification.

Three things this paper tells you that most vendor decks avoid:

1. **Week zero of any real implementation is a discovery sprint, not a build sprint**, because most MT4/MT5 white-label setups restrict Manager API access, and that blocks or degrades the Platform Activation Coach, the Early Churn Trigger and the VIP Losing-Streak Intervention before a line of code is written.
2. **WhatsApp Business Platform access is gated by Meta business verification**, which typically takes a few business days and can run to two weeks, and it must be started before Week 1, not discovered in Week 2.
3. **The Profit Truth Agent, the component most vendors lead with, is Week 7** in the correct build sequence, because it needs six weeks of clean cohort data and ad-spend integrations that take time to approve.

The sequence here reflects the order in which things actually work, not the order that makes the best slide.

The benchmark model. Every financial figure in this paper is calculated on one benchmark brokerage. The figures are a model, not an industry average: they sit at values we see often in audits, and every calculation is shown in full so you can substitute your

own numbers.

BENCHMARK ASSUMPTION	VALUE
Registrations (leads) per month	1,000
Cost per registration	$70, so $70,000 of monthly acquisition spend
Registrations that start KYC	70%, or 700
KYC completion with no intervention	40%, or 280 approved accounts (the audit range is 35 to 45%)
Approved accounts that make a first deposit	About 70%, or 200 first-time deposits: a 20% registration-to-FTD rate
Cost per funded account (CPA funded)	$350
Deposit attempts per FTD	1.2, so about 240 attempts and 40 failures a month
Net revenue contribution (NRC) per funded client	$800 over the first 180 days: the spread, commission and financing revenue attributable to the client, net of IB commission, rebates and payment costs

At these values the brokerage earns 2.3 times its acquisition spend within 180 days. Every leakage calculation scales linearly with volume.

On net revenue contribution and the book. NRC in this paper is the revenue a client generates for the brokerage through spread, commission and financing, net of IB commission, rebates and payment costs. It deliberately excludes dealing profit or loss on a B-book. A broker that internalises client flow may add realised dealing P&L to its own figure, but should not let that figure drive Expected VIP tagging or retention outreach: a model that selects the clients expected to lose most and routes them to a retention desk is the fact pattern conduct regulators have acted on. TRACE's retention components are built to run on revenue the client pays for a service, not on the client's losses.

Terms used in this paper. FTD: first-time deposit, the event that turns a registration into a funded account. CPA funded: acquisition spend divided by funded accounts. NRC: net revenue contribution, as defined above. C-ROAS: cohort return on ad spend, the NRC of a cohort of funded clients divided by the spend that acquired them, as distinct from vanity ROAS calculated on deposit volume. Speed to lead: the time from

registration to the first human contact attempt. PSP: payment service provider. IB: introducing broker. KYC: the know-your-customer identity verification a client completes before funding. Expected VIP and Proven VIP are defined in Appendix C.

What 30+ broker funnel audits found

Not because we can guess. Because we have done it. Registered as a prospect. Completed KYC. Attempted deposits. Placed trades on MT4 and MT5. Tracked every touchpoint: email, calls, WhatsApp, live chat. Waited to see if anyone followed up after signup, after a failed deposit, after KYC stalled, after the first trade. The same brokerage appears every time. The details change. The failures do not.

WHAT WE FOUND	WHAT IT ACTUALLY COSTS
No meaningful contact within the first hour of signup	The prospect cools. The CPA is spent. The conversion window closes before the rep has even opened the lead.
KYC treated as a compliance step, not a conversion gate	No intelligent nudging, no friction removal. KYC abandonment at 70 to 80% is not a compliance problem. It is a revenue problem.
Failed deposits generate no rescue playbook	Bank decline, PSP error, abandoned checkout, all treated identically: ignored. The broker writes off the CPA and buys more traffic.
Post-FTD silence	Full acquisition cost paid. The job is now to activate trading in a compliant manner. Most brokerages go quiet.
Withdrawal handled as a back-office process	Slow, opaque, unacknowledged. This is where the client relationship ends. Permanently.
Email that reads like 2010 templates	Generic, untriggered, unsequenced. The same message to a just-deposited client and a three-month-inactive ghost.

> *The most expensive executive illusion we meet in brokerage growth: "We need more leads." In most cases you do not have a lead problem. You have a leakage and allocation problem, and buying more traffic makes it worse, not better.*

What funnel leakage costs a CFD broker

Three fixable leakage points, speed to lead, KYC drop-off and failed deposits, cost the benchmark 1,000-lead-per-month brokerage roughly $600,000 to $900,000 a year. Every line below is a calculation you can run against your own numbers this afternoon, using the benchmark model above: 1,000 registrations at $70 each, 700 KYC starts, 200 FTDs, and $800 of net revenue contribution per funded client over 180 days.

Leakage point 1: speed to lead

InsideSales' 2021 Lead Response Study, across 55 million sales activities and 5.7 million inbound leads at more than 400 companies, found conversion rates were 8x higher when the first contact attempt was made within five minutes rather than between five minutes and 24 hours. 57.1% of first call attempts happened after more than a week. The two estimates in the table, the share of leads contacted late and the share of conversions lost to it, are ours, from the audits.

Monthly acquisition spend	1,000 registrations × $70 = $70,000, or $350 per funded account
Leads with no human contact inside one hour (audit observation)	~620 (62%)
Estimated conversion loss from delayed contact	~3 to 5% of total leads
FTDs lost per month from speed failure	30 to 50
Revenue cost per month at $800 NRC	$24,000 to $40,000
Annualised	**$288,000 to $480,000 from one fixable failure**

Leakage point 2: KYC drop-off

Fenergo's 2024 survey of more than 450 banking executives found 67% had lost clients to slow or inefficient onboarding, and Signicat's 2022 survey of 7,600 European consumers found 68% had abandoned a financial application during onboarding. KYC/AML friction in heavily regulated flows drives abandonment as high as 70 to 80%, on Jumio's figures. In the brokerages we audit, unassisted KYC completion runs at 35 to 45%.

Registrations reaching KYC per month	700
KYC completion rate, no intervention	40%, or 280 approved
Completion uplift from active, reason-specific nudging	+8 to 12 percentage points, or 56 to 84 additional completions
Recovered completers who go on to fund (below the 70% organic rate: a stalled applicant is less committed)	25%, or 14 to 21 additional FTDs
Revenue value per month at $800 NRC	$11,200 to $16,800
Annualised	**$134,000 to $202,000, from automations that cost weeks to build**

Leakage point 3: failed deposits

The average brokerage has no visibility into why deposits fail. Bank decline, PSP error and abandoned checkout are treated as one undifferentiated loss. Across the brokerages we audit, 15 to 20% of deposit attempts fail, and 40 to 60% of those are recoverable with a playbook that treats each failure type differently. The window for recovery is short: Oxford Economics research for Checkout.com found 45% of consumers will not retry after a single false decline and 42% never return to the site at all.

FTD attempts per month (200 FTDs at 1.2 attempts each)	~240
Failed deposits (about 17% of attempts)	~40
Recoverable with an active playbook (50%, the middle of the 40 to 60% range)	~20
Revenue recovery per month at $800 NRC	$16,000
Annualised	**$192,000, money you already spent to acquire**

Combined estimate across the three fixable leakage points: $288,000 to $480,000, plus $134,000 to $202,000, plus $192,000, so roughly $600,000 to $900,000 per year at the benchmark 1,000-lead-per-month operation. That is the equivalent of roughly nine to thirteen months of the brokerage's acquisition spend. Scale linearly for your volume. None of these require a new CRM. All of them require ownership of your event stream.

The funnel metrics broker leadership must own weekly

These are not vanity metrics. Each one maps directly to a leakage point in the funnel. If your team cannot answer these with actual numbers by end of week, you are managing stress, not growth.

PILLAR	METRICS LEADERSHIP MUST OWN WEEKLY
Truth	Attribution coverage: % of leads with persisted source and campaign, plus identity mapping integrity. Event stream health: missing-event rate and ingestion latency.
Economics	CPA funded by campaign, geo and entity. True cohort ROAS (C-ROAS) to date, plus mature-gated 7, 30, 90 and 180-day views where the cohort is ready. Forecast C-ROAS at 90 and 180 days with downside, base and upside for immature cohorts. Payback ratios at the same horizons.
Execution	Contact success rate; lead-to-call time. Weekly leakage map: top leakage stage by source and by rep.
FTD and retention	Deposit failure rate (decline, abandon, PSP error) and recovered deposit rate. Time to first login and first trade. 7-day and 30-day post-FTD active rate. 3-trading-day inactivity reactivation rate.

Monday morning: five non-negotiables

Before deploying anything in TRACE, verify you can answer these five questions with numbers, not estimates:

1. Can you trace a click or source to `lead_created` to `deposit_success` (the FTD) and tie it to net revenue contribution? If not, Truth is your immediate priority.
2. What is your current speed to lead and contact success rate by rep and by desk? If you do not know, you are not managing it.
3. What percentage of deposit failures are bank declines versus PSP errors versus abandoned checkouts? If you cannot split this, you cannot recover them.
4. What happens to a funded client in the 72 hours after their first deposit? Map it. Every gap is a revenue leak.
5. What is your IB channel's net revenue contribution after commission? If you are optimising IB volume without this number, you may be paying to shrink margins.

SECTION 1

Why smaller brokerages stall: the Operator Trap

Smaller brokerages lose their one structural advantage, speed, to the Operator Trap: leadership so consumed by daily survival that improving the funnel starts to feel like a risk.

1.1 The Operator Trap

Smaller brokerages have one structural advantage over the institutional players: speed. The ability to test, pivot and execute faster than a platform with 500 employees and three compliance layers. In practice, most lose that advantage before they lose their budget.

They get trapped in what this paper calls the Operator Trap: leadership so consumed by day-to-day survival, FTD pressure, deposit volatility, staff turnover, compliance noise, that innovation stops being a lever and starts feeling like a liability.

The evidence is not theoretical. Auditing 30+ brokerages, going through the complete prospect and client journey from website visit through KYC, deposit, trading and withdrawal, reveals the same pattern every time. The same funnel gaps. The same missing touchpoints. The same absence of intervention at moments that matter. Each audit covers:

- The website, messaging, UX, funnel structure and conversion rate optimisation
- Registering, completing KYC, attempting deposits, and placing trades on MT4/MT5
- Every touchpoint: email, calls, WhatsApp/Telegram, live chat
- Whether anyone follows up after signup, after KYC, after a failed deposit, and post-FTD
- Retention mechanisms after trading stops and during withdrawal
- Timing, tone, and whether communications are personalised or generic

After repeating this across dozens of brokers, the finding is consistent: most brokerages, particularly smaller ones, leave a large amount of revenue on the table. Not because they lack ambition. Because improvement feels like risk when you are already treading water.

Email communications read like 2010 templates. Marketing materials are recycled. CRM workflows are recycled. And the sad part: brokers usually know this. They just cannot act on it inside a reactive operating model.

> Small brokers need to act like speedboats: fast to pivot, aggressive about compounding gains. Instead, most mimic enterprise behaviours they have not yet earned the right to copy.

1.2 Reactive operations, as it shows up in real broker funnels

Once a brokerage slips into reactive operations, the funnel stops being something leadership designs and starts being something leadership chases. The clearest diagnostic is what happens, or more precisely what does not happen, at the most important moments in the client journey. It shows up in the five failure modes tabulated at the top of this paper: speed to lead treated as optional, KYC treated as compliance rather than conversion, failed deposits that trigger no rescue, silence after the first deposit, and withdrawals handled as a back-office process. Each is a moment where the funnel needed a designed intervention and got none, and each is where the acquisition cost already paid is written off.

1.3 The CRM gravity well

The Operator Trap does not exist in isolation. It is reinforced by a structural blocker inside most broker organisations: the CRM becomes the gravity well. Even when the C-suite knows the CRM is holding the business back, leadership avoids it, because every meaningful improvement feels like a multi-month project with political risk attached.

- Tracking and attribution are broken: overwritten sources, missing click IDs, disconnected from funded outcomes.
- Endpoints and webhooks are missing or non-functional, so anything beyond basic automations becomes painful and expensive.
- Support tickets sit unresolved for months. Adding a minor feature or accessing your own data becomes an ordeal.

If your CRM cannot expose events and endpoints, you do not own your funnel. Your vendor does. And if you do not own your funnel, you cannot build allocation intelligence.

1.4 The mirror worth studying: prop firms

One of the most instructive market shifts in the early-to-mid 2020s was how prop firms raised the bar on marketing, conversion optimisation and product experience, and how most retail CFD brokers missed the lesson.

Prop firms operate more like e-commerce brands: aggressive performance marketing with conversion-first landing pages, heatmaps and A/B testing as operational habits, and UX that removes friction rather than adding it. The lessons a regulated broker can take are the operational ones: conversion-first pages, testing as a habit, friction removal in onboarding, milestone systems tied to education and platform mastery, and segment-specific content that drives engagement without spamming. The incentive mechanics prop firms use, verified "best trade" leaderboards, monthly trading competitions, tiered prizes, are largely unavailable to a CFD broker, as the note below explains.

Important. Gamification, bonus structures and trading competitions are heavily restricted or prohibited under MiFID II, FCA COBS and the UAE SCA, DFSA and FSRA conduct rules: ESMA's 2018 product intervention measures restricted incentives to trade CFDs across the EU, and the FCA's permanent rules from August 2019 require firms to stop offering monetary and non-monetary inducements to encourage trading. Always obtain regulatory sign-off before implementing any incentive mechanism. See Appendix E.

SECTION 2

No tracking, no ads: attribution before paid media

2.1 The doctrine

> *If you cannot trace an ad click to a funded account and net revenue contribution, you should not be scaling paid ads.*

This is not a nuanced position. It is an operational prerequisite. Without attribution truth, performance marketing is not performance marketing. It is expensive traffic allocation with no feedback loop.

2.2 Why broken tracking is the hidden cause of so much broker dysfunction

When attribution is broken, three things happen simultaneously. Marketing gets blamed for sales problems. Sales gets blamed for funnel problems. Leadership concludes "we need more leads" when the actual problem is leakage they cannot see.

The funnel is not empty. It is leaking, and because you cannot see where, the instinct is to pour more traffic in to compensate. This is reactive operations in its most expensive form.

2.3 The speed-to-lead reality

The InsideSales study cited in the leakage calculation above makes the general case: conversion 8x higher when the first attempt lands inside five minutes, and 57.1% of first call attempts made after more than a week.

The broker version of this problem is familiar: automation fires instantly, but human contact happens days later, or never. Speed to lead is not a sales discipline problem. It is a systems problem.

2.4 What Truth actually means in a brokerage

When this paper refers to Truth, it means three distinct layers, each dependent on the one before it.

Truth Layer 1, attribution truth (click to lead). You can persist source and click identifiers through registration without losing or overwriting them. Every lead has a traceable origin.

Truth Layer 2, funnel truth (lead to funded). You can connect operational events to the same identity: contact attempts, KYC status changes, deposit attempts with failure reasons, and FTD success.

Truth Layer 3, economics truth (funded to net revenue contribution). You can tie funded cohorts back to net revenue contribution and payback windows, not to vanity metrics like registrations or deposit volume.

If any of these layers is missing, the CMO is optimising blind and the CEO is making budget decisions based on an incomplete picture of the business.

2.5 The minimum viable truth stack

CONCEPT	WHAT IT IS	WHY IT MATTERS FOR BROKERS
Click ID	Unique identifier appended to an ad click, for example GCLID from Google	Allows you to match a funded account back to the exact ad click that originated it.
Offline conversion	A conversion event that happens outside the website, inside your CRM or PSP	Tells Google or Meta: this click became an FTD. Enables campaign optimisation toward funded accounts, not registrations.
Enhanced Conversions for Leads	Google's upgrade to offline conversion import using hashed user data	Improves match rate when click IDs are lost to redirects, cross-device or cookie loss.
Conversions API (Meta)	Server-side event stream sent directly from your CRM or backend to Meta	Bypasses browser tracking limitations. Sends real milestones (KYC approved, FTD, reactivation) so Meta optimises against actual value.
Conversion adjustments	Google feature allowing you to restate a conversion value after the fact	Upload a quick FTD signal, then restate the value inside the adjustment window with the NRC realised so far or the forecast. Note the windows: Google does not import offline conversions uploaded more than 90 days after the click, 63 days for enhanced conversions for leads, and value adjustments must be made within 54 days of the original conversion. Appendix B is the checklist for this stack.

Pixels track interest. Offline and server-side events track revenue. If your stack cannot push FTD and NRC value back to the ad platforms, your media buyers will always look worse than they are, and your budget decisions will always be wrong.

2.6 Why this belongs in the C-suite, not marketing

Broken tracking does not just waste marketing budget. It creates and perpetuates the Operator Trap. When leadership cannot see the funnel clearly, they are forced to guess. When they guess, teams fight over narratives. When teams fight over narratives, the organisation stays reactive, and the solution is always the same: buy more traffic.

This is why TRACE begins with T for Truth. If you do not own the event stream, you cannot deploy routing intelligence, FTD rescue or churn engineering in a way that compounds.

SECTION 3

The broker funnel as an event stream

A funnel is an event stream: every gate produces a timestamped event with a stable ID, and without that you cannot see where revenue leaks. Most broker organisations think in departments. Marketing generates leads. Sales converts. Ops handles KYC and payments. Retention reactivates later. This departmental model is exactly how brokers end up in reactive operations, because no single team owns the chain from click to net revenue contribution.

> Your funnel is an event stream. If you cannot observe events with timestamps and stable IDs, you cannot identify leakage. And if you cannot identify leakage, you cannot engineer FTD or retention. You can only buy more traffic.

The event stream must cover both journeys: the **client journey** (click to lead to FTD to MT4/MT5 behaviour to retention outcomes to NRC) and the **IB journey** (IB to referred leads to funded outcomes to commission accrual to payout to disputes).

3.1 The funnel as gates

Treat the client journey as a sequence of measurable gates. Each gate produces an event with four attributes: **who** (lead ID / client ID), **when** (UTC timestamp), **where** (source, campaign, geo, entity, IB) and **what happened** (event type plus metadata).

Acquisition to sales. Ad click captured (click identifiers persisted to the session or user record) · lead created (registration complete, source and campaign persisted) · IB linked to lead (`ib_id` / `sub_ib_id` attached and immutable, IB channel only) · lead assigned (rep or desk ownership set, SLA clock starts) · first contact attempt (call, WhatsApp, email or chat attempt logged with timestamp) · first contact success (connected or replied, not merely attempted) · first qualification call completed (outcome logged).

Onboarding to funding. KYC started · KYC submitted (documents and selfie submitted, status becomes pending) · KYC decision (approved, rejected or needs-more-info, plus reason category) · deposit attempted (method, amount, PSP response) · deposit failed (bank decline, abandoned checkout or PSP error, plus reason code) · deposit successful, the FTD (amount, time, method).

Activation to retention. Trading account provisioned (MT4/MT5 credentials issued, server assigned) · first platform login · first trade (asset class, optional leverage band).

Post-FTD retention signals. Three consecutive trading days with no trade after the first trade (Monday to Friday only, excluding weekends and broker holidays) · retention or reactivation touchpoint logged (human or automated, channel and intent recorded).

Trust and lifecycle moments. Withdrawal requested (amount, reason if available) · withdrawal completed or failed (time to withdraw, reason codes).

IB economics moments. IB commission accrued (rule version and calculation basis) · IB commission paid (payout executed, timestamp) · IB dispute or open case logged.

> If you cannot reliably capture these events, you do not have a funnel. You have anecdotes. Appendix A is the full checklist.

3.2 Minimum event stream schema, IB-ready

Every event in the stream must carry this core schema.

FIELD	DESCRIPTION	EXAMPLE
`event_name`	Standardised event type	`deposit_failed`
`event_time_utc`	Timestamp in UTC	`2026-02-23T14:05:12Z`

FIELD	DESCRIPTION	EXAMPLE
`lead_id`	Pre-KYC identity key	`LID_928144`
`client_id`	Post-KYC identity key	`CID_482910`
`entity` / `brand`	Legal entity or brand context	Offshore Entity A
`geo`	Country or region	`AE`
`source`	Marketing source	`google_ads` / `ib`
`campaign_id`	Platform campaign identifier	`1234567890`
`click_id`	Click identifiers persisted	`gclid=…` / `fbclid=…`
`ib_id`	Referring IB identifier, if IB-sourced	`IB_1203`
`sub_ib_id`	Sub-IB identifier, if tiered	`SIB_44`
`ib_tier`	IB tier at time of event	`Tier_2`
`ib_agreement_id`	Commission plan version	`AGREEMENT_2026_01`
`rep_id`	Sales owner, if assigned	`REP_12`
`channel`	Touchpoint channel	`whatsapp` / `phone`
`metadata`	Event-specific fields	reason codes, amounts

Alongside the live fields, store `source_first_touch`, `click_id_first_touch` and `ib_id_first_touch` as immutable copies of the first values of `source`, `click_id` and `ib_id`. Downstream processes may update the live fields; nothing may ever write to the first-touch copies. Section 5 explains why.

3.3 The leakage map

A leakage map answers one question: where are we losing potential net revenue contribution, and why? Populate it weekly. It becomes your CEO and COO operating dashboard.

STAGE	LEAKAGE DEFINITION	METRIC TO TRACK	TYPICAL ROOT CAUSES	OWNER	TRACE COMPONENT
Click to lead	Clicks that never become registrations	Registration rate %, cost per lead	Landing friction, weak offer, tracking loss	Marketing	Landing-page CRO; Attribution Integrity Monitor
Lead to assigned	Leads with no owner or SLA	Time to assign, assignment rate %	No routing rules, desk overload	Sales Ops	Lead Scoring and Routing Agent
Assigned to contact attempt	Leads not contacted fast enough	Speed to lead, SLA hit %	No SLA enforcement, poor handoff	Sales Ops	SLA and Contact Cadence Enforcer
Contact attempt to success	Attempts that never connect	Contact rate %, attempts per lead	Bad time windows, no retry logic	Sales	SLA Enforcer: channel switching and retry windows
Contact to first call	Connected leads not qualified	Lead-to-call time, show rate	Calendar friction, weak follow-up	Sales	SLA Enforcer: follow-up cadence
Call to KYC started	Qualified leads that never begin KYC	KYC start %	Confusion, trust barriers, poor guidance	Ops/CS	KYC Completion Concierge
KYC started to decision	Drop-offs, rejects, stuck pending	Completion %, reject %, time in state	Document friction, unclear steps, provider delays	Ops/CS	KYC Completion Concierge
KYC approved to deposit attempt	Approved users who never try to deposit	Attempt rate %	Weak CTA, payment trust gap	Sales/Marketing	KYC Completion Concierge: approved-not-funded step
Deposit attempt to FTD	Attempts that fail or abandon	Fail %, abandon %, recovered %	Bank decline, checkout UX, no rescue	Payments/Ops	Failed Deposit Rescue

STAGE	LEAKAGE DEFINITION	METRIC TO TRACK	TYPICAL ROOT CAUSES	OWNER	TRACE COMPONENT
FTD to first login	Funded users who never log in	Login activation %	Provisioning delays, platform confusion	Ops/Retention	Platform Activation Coach
First login to first trade	Logged in but do not trade	First trade %	Waiting for setup, low confidence	Retention	Platform Activation Coach
First trade to 3-day no trade	Early inactivity window	3-day inactivity %	Waiting for setup, early losses	Retention	Early Churn Trigger
Withdrawal requested to completed	Withdrawal delays create churn risk	Withdrawal turnaround time, complaint rate	Back office delays, comms gaps	Ops/Retention	Withdrawal Trust Playbook
IB to linked lead	IB referrals lose attribution	IB attribution persistence %	Lost referral params, CRM overwrite	IT + IB Manager	IB Attribution Monitor

SECTION 4

What is TRACE? The five stages, in order

TRACE is a deployment framework of AI agents and intelligent automations that moves a brokerage from reactive firefighting to engineered conversion and retention. It does not require a CRM migration. It requires ownership of your event stream and the discipline to deploy in the right order.

Most brokerages do not have an AI problem. They have an execution sequencing problem. The typical failure pattern is to jump straight to tools, a chatbot, a few automations, a dashboard, without first fixing the one thing that makes every automation valuable: Truth.

The thesis: if you can predict behaviour, you can engineer retention instead of reacting to churn. AI is not a chatbot. It is a data-gathering and attention-allocation machine.

STAGE	WHAT IT DOES
T · Truth	Own the event stream and attribution. Build a single source of truth that connects every event in the client and IB journey to net revenue contribution, with timestamps, stable IDs and full auditability. Truth stops the internal blame loop by replacing opinions with measurable leakage and measurable economics.
R · Route	Allocation intelligence: who gets attention, and when. Eliminate pipeline leakage by controlling attention allocation. Most brokerages treat leads as interchangeable. They are not. A brokerage can have exceptional marketers and still fail if high-intent leads sit untouched in queues while reps waste time on low-quality traffic.
A · Activate	FTD engineering: remove friction where money is won or lost. Money is lost in two places, KYC stalls and deposit failures. Activation reframes both as conversion gates, not compliance steps and back-office processes. The VITAL playbooks cover the same two gates from the growth side.
C · Churn Engineer	Predict behaviour, engineer retention. Stop paying full acquisition cost for partial lifetime value. This is the stage most brokerages ignore, and the one that most directly determines whether the business compounds or stalls.
E · Economics	Run the brokerage on net revenue contribution. Most brokerages can see leads, registrations and deposits. Very few can see NRC by cohort, payback windows, true ROAS tied to funded outcomes, or IB channel profitability after commission.

TRACE is not "buy a new tool." It is a **deployment order**: a sequence of capabilities that compound on each other. Truth enables everything else: Route, Activate and Churn Engineering all run on the event stream, and Economics is only as good as the attribution beneath it. The rollout in Section 7 installs Truth in stages and pulls the two Activate automations forward to Weeks 3 and 4, ahead of routing intelligence, because they recover acquisition spend already paid and need only the deposit and KYC events wired in Week 2. Deploy without Truth and you will rebuild everything twice.

TRACE applies equally to paid acquisition funnels and IB partner networks. An IB channel without durable attribution and payout auditability is not a growth channel. It is a commission expense. The framework must produce two distinct truths: client journey truth and IB journey truth.

SECTION 5

The 16 TRACE components: 3 AI agents and 13 automations

TRACE is 16 components, three AI agents and 13 deterministic automations, each firing on an event or a schedule and each moving one primary KPI. Brokerages do not lose because they lack ideas. They lose because they cannot reliably execute improvements inside a messy, vendor-locked stack. TRACE delivers through a combination of **AI agents** (predictive, scoring-based components that learn from data) and **intelligent automations** (rules-based trigger systems that execute defined playbooks). Both are essential.

Design principles. Every automation is event-triggered, with a clear "when" based on the event stream, or scheduled where the signal is a daily or weekly aggregate. Every component moves one primary KPI, listed first in its specification, measurable, owned and reported weekly, with supporting metrics beside it. Every component is deployable using standard broker integration surfaces: CRM webhooks, PSP callbacks, MT4/MT5 sync, ad platform imports, IB ledger. Automations ship first, and the AI agents ship as rules-based first versions; models replace the rules once you have clean labels and stable truth. The "Built in" column maps each component to its week in the Section 7 rollout.

Executive map

STAGE	COMPONENT	PRIMARY KPI	TRIGGER	OWNER	BUILT IN
T	Attribution Integrity Monitor	CPA-funded accuracy	`lead_created`; daily audit	IT + CMO	Week 1 (field locks), Week 7 (alerts)
T	IB Attribution Monitor	IB payout accuracy; dispute rate	`ib_linked_to_lead`; `lead_created`; daily audit	IT + IB Manager	Week 1 (fields), Week 11 (audit)
T	Profit Truth Agent *(AI)*	CPA funded; NRC; true and forecast C-ROAS	`deposit_success`; weekly rollups	CMO + CEO/COO	Week 7
T	Event Stream Monitor	Pipeline visibility	Continuous; nightly	IT / Data	Weeks 1 to 5

STAGE	COMPONENT	PRIMARY KPI	TRIGGER	OWNER	BUILT IN
R	Lead Scoring and Routing Agent *(AI)*	FTD %; pipeline leakage	`lead_created`; `lead_updated`	Head of Sales + Sales Ops	Week 6 (rules), Week 12 (re-tiering)
R	IB Lead Router	IB FTD %; partner satisfaction	`ib_linked_to_lead`	Head of Sales + IB Manager	Week 6
R	SLA and Contact Cadence Enforcer	Speed to lead; contact rate	`lead_assigned`; SLA breach	Sales Ops	Week 1 (alert), Week 6 (tiered ladder)
A	KYC Completion Concierge	KYC completion rate	`kyc_started`; `kyc_pending`; `kyc_rejected`	Ops / CS	Week 4
A	Failed Deposit Rescue	FTD %	`deposit_failed` with reason code	Payments/Ops + Sales	Week 3
A	Platform Activation Coach	FTD-to-trade activation	provisioned; no login or trade	Retention / CS	Week 8
C	Expected VIP Tagger *(AI)*	Retention ROI	Early post-FTD window	CEO/COO + Retention	Week 0 (rules), Week 9 (model)
C	Early Churn Trigger	Reactivation; retention	3 trading days with no trade after the first	Retention	Week 8
C	VIP Losing-Streak Intervention	VIP 90-day retention rate	Losing days plus drawdown	VIP Desk / Retention	Week 10
C	Withdrawal Trust Playbook	Withdrawal turnaround time; post-withdrawal retention	`withdrawal_requested`	Ops + Retention	Week 10
C	IB Partner Retention Assist	Reactivation in IB cohorts	Inactivity; withdrawal risk	Retention + IB Manager	Week 11
E	IB Economics Engine	IB channel profitability	`commission_accrued`; weekly rollups	CEO/Finance + IB Manager	Week 11

AI agents versus automations, and why the distinction matters

Only three components require machine learning: the **Profit Truth Agent** (forecast cohort ROAS modelling), the **Lead Scoring and Routing Agent** (upgrades from rules to ML scoring once it has four to six months of labelled data) and the **Expected VIP Tagger** (a behavioural prediction model, recalibrated quarterly). All three ship first as rules-based versions. Every other component is a deterministic automation: it fires on an event, executes a rule and produces an output. No model. No training data. Deployable in weeks, not quarters.

We call out the distinction because you should not buy AI you do not need. The other 13 are not less valuable for being automations. A well-engineered automation that responds to a failed deposit within minutes, or enforces speed-to-lead discipline across every rep every day without human oversight, is operational leverage most brokerages have never had. The label matters for honesty, not for diminishing the ROI.

The sequencing is deliberate: automations ship first because they generate the clean, labelled data that AI agents require. You cannot train a lead scoring model on six weeks of dirty data.

T · Truth layer

Automation 1: Attribution Integrity Monitor

Monitors whether you can reliably connect spend to leads to funded outcomes. It detects when attribution breaks: lost click IDs, overwritten sources, spikes in unknown source, and sudden drops in match rate. Attribution breakage that goes undetected for even a week means your media buyers are optimising toward the wrong outcomes.

Primary KPI	CPA funded accuracy; ROAS measurement integrity
Trigger events	`lead_created` (continuous); scheduled daily audit
Inputs required	Click ID presence, source and campaign fields, first-touch versus last-touch fields, `lead_id` to `client_id` mapping, redirect and landing metadata
Owner	IT (fix) + CMO (decision-making)

Success metrics	Stable click capture rate, stable source persistence, low unknown-source rate, stable match rate to funded outcomes
Outputs	Daily health report, incident log, suspected root-cause category

Decision logic. Alert if click ID capture rate drops below threshold; if unknown source rate spikes above threshold; if the source field is overwritten after registration, KYC or deposit; if the mismatch rate between ad platform data and CRM data rises. **Actions:** notify IT and CMO, flag affected campaigns, open an attribution incident record. **Failure modes:** redirects stripping URL parameters, cookie loss, CRM overwrite on re-registration, duplicate records, missing offline conversion feedback.

> Store immutable `source_first_touch` and `click_id_first_touch` fields. These must never be overwritten by any downstream process.

Automation 2: IB Attribution Monitor

IB growth is only real if IB attribution is durable and auditable. This prevents the most common IB failure mode: an IB claims a client, your systems cannot prove the attribution path, and disputes become routine. It detects lost `ib_id`, attribution overwrites and duplicate client-to-IB mapping issues before they become payout problems.

Primary KPI	IB payout accuracy; IB dispute rate reduction
Trigger events	`ib_linked_to_lead`; `lead_created`; scheduled daily audit
Inputs required	`ib_id`, `sub_ib_id`, `ib_tier`, `ib_agreement_id`, immutable first-touch attribution fields, identity mapping
Owner	IT + IB Manager
Success metrics	High IB attribution persistence, low dispute rate, clean lead-to-client-to-IB mapping
Outputs	IB attribution health report, incident log, root-cause category

Decision logic. Alert if an IB-linked lead loses its `ib_id` at any subsequent stage; if one `client_id` maps to multiple IBs; if IB traffic landing paths fail to attach `ib_id` consistently. **Actions:** notify IT and the IB manager, quarantine the affected cohort for payout review before commission runs.

Store `ib_id_first_touch` as immutable. Version all commission agreements using `ib_agreement_id`. This is the audit trail for every future dispute.

AI Agent 1: Profit Truth Agent

Most broker teams do not have a ROAS number. They have a debate. Marketing shows registrations. Sales argues lead quality. Finance sees different numbers. The Profit Truth Agent ends that debate by producing a single source of economic truth: CPA funded by campaign, geo and entity; true cohort ROAS (C-ROAS) tied to net revenue contribution; and forecast C-ROAS for immature cohorts. Appendix D works one cohort through the method end to end. It is the instrument panel that tells the CEO whether to scale, hold or kill a channel.

Primary KPI	CPA funded; true C-ROAS; forecast C-ROAS; payback ratios
Trigger events	`deposit_success`; weekly NRC rollups; cohort maturity thresholds
Inputs required	`click_id`, source and campaign, FTD amount, NRC by `client_id`, IB commission data, cohort assignment dates
Owner	CMO + CEO/COO
Success metrics	All funded cohorts tracked to NRC, mature and immature cohorts clearly separated, a weekly decision output produced
Outputs	Weekly C-ROAS report, the scale, hold or kill memo, payback dashboard by cohort

Decision logic. Calculate CPA funded as spend divided by funded accounts, by campaign, geo and entity. Calculate true C-ROAS on mature cohorts only, flagging immature cohorts explicitly. Produce forecast C-ROAS for immature cohorts with downside, base and upside ranges. Calculate payback ratios at 7, 30, 90 and 180 days, mature-gated. Flag channels where the payback ratio is below threshold or forecast C-ROAS is below break-even. **Actions:** produce a weekly scale, hold or kill memo; surface to CEO and CMO; flag channels requiring a decision within 48 hours.

True C-ROAS on immature cohorts is misleading. Always flag maturity status prominently. A channel that looks unprofitable at 30 days may be highly profitable at 90, and vice versa. Never make scale or kill decisions on immature cohorts without forecast C-ROAS.

Automation 3: Event Stream Monitor

The infrastructure watchdog for TRACE. It monitors the event pipeline for missing events, ingestion latency and identity resolution failures. An attribution system that worked last Tuesday may be silently broken today, and without this automation you find out when your weekly numbers make no sense. Data quality is not an IT concern. It is a revenue concern.

Primary KPI	Pipeline visibility; missing-event rate; ingestion latency
Trigger events	Continuous monitoring; nightly batch reconciliation
Owner	IT / Data
Success metrics	Under 1% missing event rate, under 5 minutes ingestion latency, over 99% identity mapping completeness

Decision logic. Alert if event counts for any `event_name` drop more than 20% against the same period the prior week; if ingestion latency exceeds five minutes for any critical event type; if `lead_id` to `client_id` mapping completeness drops below threshold. Nightly, reconcile CRM record counts against event stream counts for key milestones. **Failure modes:** CRM webhook failures, PSP callback drops, MT4/MT5 sync gaps, network timeouts, schema changes breaking parsers.

> Instrument this before deploying any downstream agent. A routing agent fed by a broken event stream will route incorrectly, and you will not know until you audit the outputs.

R · Route layer

> Speed to lead is decided by the system, not by the rep: the queue, the SLA timer and the escalation ladder determine who is called inside five minutes.

AI Agent 2: Lead Scoring and Routing Agent

Most brokerages route leads by round-robin or first-available. This replaces gut feel with a signal-based priority queue: each lead arrives with a score based on source quality, geo, behavioural signals and prior engagement, and is routed to the right rep with an SLA timer started immediately.

Primary KPI	FTD % by routed segment; cost per funded account; SLA hit rate
Trigger events	`lead_created`; `lead_updated` (score revision on new signals)
Inputs required	Source, geo, campaign quality score, device type, time of registration, previous engagement, desk capacity
Owner	Head of Sales + Sales Ops
Success metrics	SLA hit rate, FTD rate by routed tier, cost per funded account by tier, rep workload balance
Outputs	Priority queue by score, assignment records with SLA timestamps, daily routing efficiency report

Decision logic. Score each lead on source quality, geo historical FTD %, time-of-day intent signal, device and campaign tier. Assign to desk or rep based on score, current capacity and language/timezone match. Start the SLA timer immediately on assignment. Revise the score on new signals within the first 30 minutes, such as email verified or phone confirmed. **Failure modes:** stale scoring models not updated with recent geo performance, desk capacity not reflected in real time, rules overriding scores arbitrarily.

> Start with rules-based scoring, three to five signals, and re-tier from the data every quarter. Add ML scoring once you have four to six months of clean labelled data and at least 200 labelled conversions (`lead_id` plus FTD outcome).

Automation 4: IB Lead Router

IB-sourced leads have fundamentally different dynamics than paid traffic. The referred client has a prior relationship with the IB, which means the quality and urgency of handling affects not just this client's conversion but the IB's willingness to keep sending volume.

Primary KPI	IB FTD %; IB partner satisfaction measured by referral volume trends
Trigger events	`ib_linked_to_lead`; `lead_created` (IB-sourced)
Owner	Head of Sales + IB Manager
Success metrics	IB FTD rate against non-IB FTD rate by comparable source and geo, IB partner referral volume trend
Outputs	IB-tagged priority queue, IB relationship manager notification log, IB cohort FTD rate by tier

Decision logic. Route IB leads based on tier SLA, so Tier 1 IBs get a faster SLA than Tier 3. Match to rep by language and timezone plus IB relationship familiarity. Notify the IB relationship manager when Tier 1 or 2 IB leads arrive. Track IB-specific FTD % separately to feed the IB quality score model.

> When a good IB sends a referred client who is handled slowly, the IB notices before you do. Tier 1 IB leads should receive your fastest SLA.

Automation 5: SLA and Contact Cadence Enforcer

The enforcement layer for everything the routing agent sets up. SLA timers are only valuable if violations trigger action. This monitors every assigned lead for SLA breach, fires escalation tasks to supervisors, and initiates channel-switching when contact is not achieved, automatically and in real time.

Primary KPI	Speed-to-lead SLA hit rate; first contact success rate; lead-to-call time
Trigger events	`lead_assigned`; first contact attempt, or the absence of one; SLA threshold breaches
Owner	Sales Ops
Success metrics	Share of leads contacted within SLA, first contact success rate, average time to first contact by tier
Outputs	Real-time SLA breach log, weekly SLA performance by rep and desk, channel-switch effectiveness report

Decision logic. The ladder is tiered to the lead score. Tier A, five-minute SLA: at T+5 minutes with no attempt, alert the rep and create a high-priority task; at T+15 minutes with no success, switch channel, phone to WhatsApp or WhatsApp to email, and alert the supervisor; at T+60 minutes, escalate to the desk head and flag for manual override; at T+24 hours, move to the nurture queue and flag for the weekly leakage review. Tier B, one-hour SLA: the same ladder at T+60 minutes, T+4 hours and T+24 hours. Tier C: a same-business-day attempt, then the nurture queue.

> Do not set a single SLA for all lead tiers. Tier A leads get a five-minute SLA, the window the response-time evidence supports. Tier B can tolerate an hour. Tier C gets a same-business-day attempt, then self-serve and nurture. Tier your SLAs to your scoring model.

A · Activate layer

Automation 6: KYC Completion Concierge

KYC abandonment is not a compliance inevitability. It is a revenue leak with predictable causes: document confusion, unclear next steps, rejection without guidance, and stalling at pending status with no communication. This treats the KYC journey as a conversion funnel with specific intervention points at each failure stage.

Primary KPI	KYC completion rate; time to KYC approved; pre-FTD leakage reduction
Trigger events	`kyc_started`; `kyc_pending` more than 48h; `kyc_rejected` or needs-more-info; `kyc_approved` with no deposit attempt; time-since-started thresholds
Owner	Ops / CS
Success metrics	KYC completion rate, time to approval, rejection rate by reason, human escalation rate, approved-to-deposit-attempt rate
Outputs	KYC completion rate by stage, stall reason analytics, human escalation queue, rejection reason distribution

Decision logic. On `kyc_started`, send a step-by-step document guide for the client's geo. At T+2 hours still not submitted, nudge with a specific document checklist. At T+24 hours, "we are holding your account open" with a direct link back. At T+48 hours, create a human task for CS outreach and send no further automated messages. On `kyc_rejected` or needs-more-info, send a specific correction guide based on the reason code within 30 minutes. On `kyc_pending` beyond 48 hours, alert Ops to investigate a provider-side delay. On `kyc_approved` with no `deposit_attempted` within 24 hours, send the funding guide with the payment methods for the client's geo; at 72 hours, create a sales task. Approved-but-unfunded is the gate between this component and the Failed Deposit Rescue, and nothing else covers it.

> "The time to apply (21%), the amount of personal information required (21%) and changing their mind (21%) are the main reasons applications are abandoned."
> Signicat, The Battle to Onboard 2022, a survey of 7,600 European consumers

> In our audits, the highest-ROI moment in KYC is immediately after rejection. A specific, helpful correction message within 30 minutes recovers a large proportion of rejections. Generic resubmit messages recover almost none.

Automation 7: Failed Deposit Rescue

A failed deposit is a funded account that almost happened. The client has cleared KYC, expressed payment intent, and failed at the last gate. In the brokerages we audit, this is the highest-leverage recovery point in the funnel, and most ignore it entirely.

Primary KPI	FTD % (deposit rescue contribution); recovered deposit rate
Trigger events	`deposit_failed` with reason code
Inputs required	Failure reason code, deposit amount, payment method attempted, Expected VIP flag (or the attempted amount, before an FTD exists), previous attempt history
Owner	Payments/Ops + Sales
Success metrics	Recovered deposit rate by failure type, rescue message open and response rate, time to rescue message
Outputs	Rescue message delivery log, recovery rate by failure type, alternative-method conversion, VIP escalation log

Decision logic. *Bank decline:* immediate message with two or three alternative payment methods specific to the client's geo; avoid a resubmit-same-card prompt. *Abandoned checkout:* re-engagement as soon as abandonment is detected, 30 minutes after `deposit_attempted` with no outcome, with session context and in a soft recovery tone. *PSP error:* immediate acknowledgment, manual retry instruction, and an ETA for resolution. *High expected-value client, any failure:* create an immediate VIP desk human task regardless of failure type; before a first deposit exists, the attempted amount is tested against the same Expected VIP threshold. **Failure modes:** failure reason codes not captured from the PSP, VIP threshold not configured, rescue messages firing after 60+ minutes, the same message for all failure types.

> Bank decline and abandoned checkout require completely different messaging. A bank decline needs an alternative method. An abandoned checkout needs reassurance and a frictionless return path. Never send the same rescue message to both.

Automation 8: Platform Activation Coach

A client who funds and never logs in is a CPA converted to zero lifetime value. The gap between FTD and first trade is a measurable failure window, and it is preventable. This guides newly funded clients through provisioning to first trade: platform download, login, navigating to their first instrument, and placing a trade.

Primary KPI	FTD to first-trade rate; first-login activation rate; time to first trade
Trigger events	`deposit_success` to provisioning; no login at T+48h; first login with no trade at T+72h
Owner	Retention / CS
Success metrics	Login activation rate within 48 hours, first-trade rate within 72 hours of login, time-to-first-trade distribution
Outputs	Activation funnel metrics (provisioned to login to first trade), message engagement rates, human escalation queue

Decision logic. On provisioning, send immediate confirmation plus a platform download link and a login guide specific to device and OS. At T+48 hours with no login, resend login instructions with a direct support option. At T+72 hours after first login with no trade, send a first-trade walkthrough for the relevant instrument category. The third touch is human: a CS task for a personal activation call. Two automated touches after the confirmation, no more. Clients who fund and never trade belong here; clients who trade and then stop belong to the Early Churn Trigger.

> The most common reason for no first trade is not disinterest. It is waiting for a market setup, or uncertainty about platform mechanics. The activation coach addresses both. Do not treat this as a sales follow-up. It is educational and supportive.

C · Churn engineering layer

AI Agent 3: Expected VIP Tagger

Not all funded clients are equal, and knowing which ones are likely to become high-value before they have demonstrated it gives retention teams a critical head start. This applies a scoring model in the early post-FTD window to identify clients whose deposit size, geo, source and behavioural signals suggest above-average lifetime value.

Primary KPI	Retention ROI; VIP identification accuracy at 30 and 90 days
Trigger events	T+24h after first trade, or T+48h after FTD, whichever comes first
Owner	CEO/COO + Retention
Success metrics	Tag precision at 30 days (the share of Expected VIPs tracking toward Proven), tag recall (the share of Proven VIPs who were tagged)
Outputs	Expected VIP tags with confidence scores, VIP desk routing queue, model accuracy report at 30 and 90 days

Decision logic. Score on deposit amount quartile, geo tier, source quality and early platform engagement depth. Tag as Expected VIP if the score exceeds threshold, calibrated against the 90-day realised VIP rate. Route to the VIP desk queue immediately and set VIP treatment flags in the CRM. Review tag accuracy at 30 and 90 days; recalibrate quarterly.

Start with a rules-based flag on FTD amount alone, the top 15 to 20% of historic deposit sizes (Week 0). Upgrade to a multi-signal scored model once you have about eight weeks of FTD data and four weeks of first-login data (Week 9), and recalibrate it against realised NRC as cohorts reach 180 days. Appendix C defines Expected and Proven VIP. Precision matters more than recall: flooding the VIP desk with non-VIPs destroys desk capacity.

Automation 9: Early Churn Trigger

Three consecutive trading days with no trade after the first trade is the earliest reliable churn signal we have found in retail CFD. At this stage the client has not rage-quit or rage-withdrawn. They are waiting: for a setup, for confidence, for a reason to come back. The intervention must be light. No aggressive sales. No deposit pressure.

Primary KPI	Reactivation rate within 7 days of the intervention; 30-day active rate post-intervention
Trigger events	Three consecutive trading days (Monday to Friday, excluding broker holidays) with no trade after the first trade
Owner	Retention
Success metrics	Reactivation rate within 7 days of intervention, 30-day active rate post-intervention, intervention-to-withdrawal rate

Outputs	3-day inactivity trigger log, intervention outcome report (reactivated, still inactive, withdrawn), channel effectiveness comparison

Decision logic. Count only live trading days. On trigger, fire a light-touch message: a market brief, an instrument-specific setup alert, or a platform tip. If Expected VIP, create a human task for the VIP desk the same day. If a standard client, send the automated message and create a human task if there is no response within 24 hours. **Failure modes:** weekend days counted in the 3-day window, creating false positives; message tone too salesy, which increases churn rather than reducing it; intervention not tracked to outcome.

> The message at this stage must feel like a thoughtful service touch, not a sales call. A market brief or setup alert outperforms any retention offer at this window. The client does not know they are in a churn risk protocol.

Automation 10: VIP Losing-Streak Intervention

A VIP client in significant drawdown is a predictable churn risk. The psychology of loss in trading follows well-documented patterns, and the window for a well-timed, empathetic touch is narrow. This monitors drawdown thresholds for high-value clients and triggers a human escalation before the client has disengaged, so that the conversation is a support call about leverage, risk and the client's own objectives, not a save call.

Primary KPI	VIP 90-day retention rate; share of flagged clients who receive a documented support contact; VIP churn rate
Trigger events	Losing-streak threshold met, configurable per VIP tier and calibrated to the client's leverage and trade frequency; the benchmark rule is three consecutive losing days and a drawdown beyond 20% of account equity at day start
Owner	VIP Desk / Retention
Success metrics	VIP churn rate before and after implementation, contact-to-retention rate by intervention type, post-contact 30-day retention among flagged clients
Outputs	VIP churn risk flags, VIP desk task queue, support contact outcome log, quarterly retention ROI by intervention type

Decision logic. Monitor VIP-flagged clients only. On trigger, create a human task for the assigned VIP desk rep, within two hours for Proven VIPs and the same day for Expected VIPs, with full client context: recent trading activity, P&L summary, last touchpoint, and

a recommended support approach based on the client profile. **This intervention is human-only.** Do not fire automated messages on this trigger.

> The rep contacting a VIP during a losing streak must know the context before they dial. Calling to check in without knowing the client has just had three losing days is worse than not calling at all. The script is a vulnerability and appropriateness check, written with compliance, recorded, and never tied to a deposit or a bonus.

Automation 11: Withdrawal Trust Playbook

A withdrawal request is not the end of the relationship. It is a decision point. Handled well, it is the moment that builds the most durable trust. Handled badly, with delays, unclear status and no communication, it generates the most damaging word of mouth and the most certain churn.

Primary KPI	Withdrawal turnaround time; post-withdrawal retention rate; complaint rate
Trigger events	`withdrawal_requested`; status changes; withdrawal pending more than 24h
Owner	Ops + Retention
Success metrics	Withdrawal turnaround time, average and 90th percentile; complaint rate per 100 withdrawals; post-withdrawal 30-day retention rate
Outputs	Withdrawal processing turnaround report, communication log, complaint rate by delay band, post-withdrawal retention rate

Decision logic. On request, send an automated acknowledgment within one hour with a realistic, specific processing timeline ("2 to 3 business days", not "soon"). At T+24 hours still pending, send a status update with a proactive explanation if a delay is expected. At T+48 hours, escalate to Ops with a retention notification. On completion, send confirmation and a thank-you; for clients who withdrew less than 70% of their equity, an optional re-engagement message seven days later; for VIP clients, a personal call from the retention desk.

> Post-withdrawal re-engagement is the most underused retention lever we see in retail CFD. A client who withdrew and had an excellent experience is your best source of re-deposit and referral. Treat withdrawal completion as the start of a re-engagement sequence, not the end of the relationship.

Automation 12: IB Partner Retention Assist

When an IB's referred clients go inactive or withdraw at elevated rates, the IB notices before you do. If the IB experiences this without acknowledgment or coordinated response, they move volume to a competitor. This creates a coordinated save workflow that loops in both the retention team and the IB relationship manager, within defined communication boundaries.

Primary KPI	IB cohort reactivation rate; IB referral volume trend; IB-sourced 90-day retention
Trigger events	Inactivity concentration in an IB-sourced cohort; withdrawal rate spike; IB-specific 3-day inactivity cluster
Owner	Retention + IB Manager
Success metrics	IB cohort reactivation rate, IB referral volume trend post-intervention, logged IB-coordinated touch outcomes
Outputs	IB cohort reactivation analytics, IB relationship manager activity log, coordinated touch outcome report

Decision logic. Monitor inactivity and withdrawal rates by `ib_id` cohort weekly. On threshold breach, alert the retention team and IB relationship manager simultaneously. Coordinate save tasks within policy: every IB-coordinated touch must be logged with `ib_id`, touch type and outcome. The IB relationship manager may notify the IB of the retention effort but **may not direct the IB to make sales contact with clients.** That is a compliance boundary.

> IBs must not become a shadow CRM. Every IB-coordinated client touch must be logged, policy-bound and auditable. Retention coordination is permitted; IB-directed sales contact with existing clients is not, unless your compliance function has approved a written policy for it.

E · Economics layer

Automation 13: IB Economics Engine

IB commissions are often the largest single variable cost line in a brokerage P&L, and the least interrogated (the unit economics of a retail CFD trade gives the engine-level view of the same P&L). This produces weekly profitability analysis by IB and by tier: gross

revenue contribution from IB-sourced clients, minus commissions paid, equals net IB contribution. Combined with cohort quality metrics, it drives structured tier reviews: which IBs to reward, which to restructure, and which to offboard.

Primary KPI	IB channel profitability; net IB contribution by tier; commission-to-NRC ratio
Trigger events	`commission_accrued`; weekly NRC rollups; quarterly tier review cycle
Owner	CEO/Finance + IB Manager
Success metrics	Net IB contribution by tier, commission-to-NRC ratio, IB cohort quality score, tier review action rate
Outputs	Weekly IB P&L by IB and tier, quarterly tier change recommendations, commission efficiency dashboard, negative-margin IB flag list

Decision logic. Calculate net IB contribution as NRC from the IB cohort minus commissions paid, by `ib_id` and tier. Calculate the commission efficiency ratio as commissions divided by NRC, flagging IBs above threshold. Track cohort quality metrics by IB to separate volume IBs from value IBs. Produce quarterly tier promotion and demotion recommendations based on net contribution and cohort quality.

Rule versioning is non-negotiable. You need to prove which plan applied when, in any dispute resolution. Every commission calculation must reference a specific agreement version (`ib_agreement_id`).

These are not AI features. They are a compounding system of AI agents and intelligent automations. Truth creates visibility. Route allocates attention. Activate removes deposit friction. Churn Engineering intervenes in predictable post-FTD risk windows. Economics forces every decision against net revenue contribution. The automations execute your rules at scale; the AI agents learn from your data and improve over time.

SECTION 6

How to deploy TRACE on your existing CRM

You do not need a new CRM to deploy TRACE. You need ownership of your event stream and the ability to trigger actions from it. This section explains how to deploy in the real world, even if your CRM is legacy, vendor-locked or missing endpoints.

6.1 The three paths

Path A, expose the events (fastest). Use when your CRM has usable APIs or webhooks, or your vendor will implement them. The goal is to stream key events out of the CRM into your agent layer, and push actions back in: tasks, tags, routing. What you must ask for: `lead_created` with source and campaign persistence; `lead_assigned` with rep ownership and SLA start; contact attempts and outcomes; KYC state changes with reasons; deposit attempted, failed and successful with reason codes; withdrawal requested and completed; IB linkage fields if you run an IB channel.

Path B, build a truth layer (best long term). Use when CRM endpoints are weak, or you do not trust CRM data consistency. Build a broker-owned event stream and warehouse fed by every system that matters: CRM, PSP and payments, MT4/MT5, ad platforms for spend and conversion feedback, and the KYC provider. Your agents and dashboards operate on the truth layer, not on whatever reporting the CRM gives you. This solves vendor lock without forcing an immediate migration.

Path C, migrate (only when the vendor blocks ownership). Use when your CRM vendor refuses endpoints, restricts data access, or makes integration impossible. Migrate only after you have proven the ROI and clarified the exact event stream you need. The mistake brokers make is migrating without a clear truth model. TRACE prevents that by giving you the exact requirements first.

6.2 The minimum viable integration surface

You do not need every integration under the sun. You need the few surfaces that control the funnel:

1. **CRM or lead database:** lead and client identity mapping, rep assignment and tasks, contact logs
2. **KYC provider:** started, submitted, approved, rejected, with reason codes where possible
3. **PSP / payments:** deposit attempted; deposit failed with reason (bank decline, abandoned checkout, PSP error); deposit success
4. **Trading platform (MT4/MT5):** account provisioned, first login, first trade, basic activity metrics

5. **Ad platforms (Google/Meta):** spend by campaign, adset and ad; conversion feedback loop via offline conversions or server-side events

If any one of these is missing, TRACE still works, but your ceiling drops sharply.

6.3 The minimum viable event list

These are the event names used throughout this paper and in Appendix A. Whatever your CRM, PSP or platform calls them internally, map them to these so that every automation, dashboard and dispute speaks one vocabulary.

- `lead_created`: registration complete; source, campaign and click identifiers persisted, first-touch copies locked
- `lead_updated`: a score-relevant field changed, such as email verified or phone confirmed
- `ib_linked_to_lead`: `ib_id` and `sub_ib_id` attached and immutable (IB channel only)
- `lead_assigned`: rep or desk ownership set; the SLA clock starts
- `first_contact_attempt`: channel, timestamp and rep
- `first_contact_success`: connected or replied, not merely attempted
- `qualification_call_completed`: outcome logged
- `kyc_started`, `kyc_submitted`, `kyc_decision`: the decision carries approved, rejected or needs-more-info plus a reason category; `kyc_approved`, `kyc_rejected` and `kyc_pending` are the states the automations key on
- `deposit_attempted`: method, amount, currency and PSP
- `deposit_failed`: failure category (bank decline, abandoned checkout or PSP error), reason code and retry count
- `deposit_success`: amount, method and time; the first one is the FTD
- `account_provisioned`: MT4/MT5 credentials issued, server assigned
- `first_login`: first platform session
- `first_trade`: asset class or symbol, optional leverage band

- `three_trading_days_no_trade`: emitted by the daily calendar job when a client who has traded goes three consecutive trading days without a trade
- `withdrawal_requested`, `withdrawal_completed`, `withdrawal_failed`: amount, time to complete and reason codes
- `commission_accrued`, `commission_paid`: rule version (`ib_agreement_id`), basis and amount (IB channel only)
- `ib_dispute_opened`: an IB dispute or open case logged against a client mapping

6.4 Ownership model

TRACE does not run on goodwill. Without clear ownership, the system will be built, then slowly neglected.

- **Event stream:** IT/Data own the pipeline, schema and data quality monitoring
- **Attribution truth:** CMO + IT own the chain from click to NRC
- **Automation logic:** Marketing and Sales Ops own the rules, thresholds and SLA configurations
- **Component outputs:** each component's designated owner, as specified in Section 5
- **Weekly TRACE review:** CEO/COO, the leadership layer that enforces scale, hold or kill discipline

6.5 Practical sequencing

Truth first: implement the minimum viable event stream in stages, CRM, then PSP and KYC, then MT4/MT5, because without it nothing else is reliable. Then the two Activate automations that recover spend already paid, failed deposit rescue and the KYC concierge, which need only the deposit and KYC events. Then Route: lead scoring and routing, tiered SLA enforcement, and the weekly leakage map cadence. Then Churn Engineering: post-FTD activation and inactivity, the VIP losing-streak intervention, the withdrawal trust playbook. Economics becomes actionable once six weeks of cohort data exist, which is why the Profit Truth Agent is Week 7, and the IB layer follows in Week 11.

SECTION 7

The 90-day TRACE rollout plan

TRACE deploys in 12 build weeks after a mandatory Week 0 of prerequisites: about 90 days of build time, in three four-week phases, with the first revenue recovery inside 30 days. Elapsed time is usually 13 to 16 weeks, because the vendor approvals in Week 0 (MT4/MT5 Manager API access, WhatsApp business verification, ad-platform API access) take one to three weeks to come back and should be requested before the build team is booked. This plan is built for small-to-mid brokerages: a team of 3 to 8 people, one developer (or a shared one), a CRM admin, and a Head of Sales or COO who owns the outcome. It is not built for enterprise. It is built for the team that needs to see revenue impact in the first 30 days or it will lose the mandate to continue.

> This is **12 build weeks plus Week 0, not a 12-week calendar from a standing start.** A team that skips Week 0, or tries to run the three phases in parallel, will ship nothing properly. It is structured as three four-week phases plus a mandatory pre-work block that must complete before Week 1. A small team that completes Phase 1 fully will recover more revenue than a larger team that half-ships all three phases simultaneously.

Week 0: technical prerequisites

Week 0 is not optional. Every item is a dependency that blocks downstream builds. Discovering these in Week 3 means rebuilding. Discovering them in Week 0 means planning. This block takes 3 to 5 working days of effort across the CRM admin, developer, and Head of Sales or COO, but its vendor requests take one to three weeks to come back, so start it before any build code is written.

0A. CRM webhook audit. Verify that your CRM can fire outbound webhooks on the specific events TRACE needs. Many CRMs support webhooks on record creation but not on field changes, and TRACE needs field changes: lead source being overwritten, KYC status updating, rep assignment changing. Check: can the CRM fire a webhook when `lead_source` is updated on an existing record? On rep assignment change? Does it expose a stable, immutable `lead_id` that persists through KYC and deposit? Can the original lead source field be locked after first write? If no to any, escalate to the vendor before Week 1 and budget an extra 3 to 5 days.

0B. WhatsApp Business Platform access. *Critical path.* Access runs through a Business Solution Provider and Meta business verification, which typically takes a few business days and can run to two weeks; until it completes you are capped at 250 business-initiated conversations a day. Meta says template review can take up to 24 hours; budget 24 to 48 hours per template, and templates can be rejected. Sign up with a Business Solution Provider, submit business verification documents to Meta, and draft the templates needed for Weeks 1 to 4 now: SLA breach alert, deposit failure rescue (three variants), KYC nudge (three variants). If WhatsApp is not approved by Week 2, fall back to SMS for the SLA enforcer and email for deposit rescue. Build both channels so you are not blocked.

0C. MT4/MT5 Manager API access check. *Critical path.* This is the single most common blocker in TRACE implementations. White-label brokers often do not have direct Manager API access; it is controlled by their provider. Confirm you can pull account provisioning events, first login timestamp, first trade timestamp, daily P&L per account and current account equity. If your MT4/MT5 is hosted by a white-label provider, request API access in writing this week, because granting it can take 1 to 3 weeks of vendor negotiation. Without access, the Week 5 platform events must be polled from daily reports, Week 8's Activation Coach and Early Churn Trigger run as a once-daily batch off the report import, and Week 10's Losing-Streak Intervention cannot be built as specified.

0D. PSP failure code mapping. Week 3 builds three separate rescue playbooks for bank decline, PSP error and abandoned checkout. This only works if your PSP returns structured, machine-readable failure codes. Many return free text. List every PSP you use, since each has a different schema. Build a mapping table now: PSP code to TRACE category. If a PSP returns free text, you need a regex or keyword classifier. Budget half a day to a day per additional PSP beyond the first.

0E. Ad spend API setup. *Critical path for Week 7.* The Profit Truth Agent requires daily spend from the Google Ads API and Meta Marketing API, joined to your event stream by campaign ID. Google requires a developer token: Basic access needs Google Cloud brand verification first and is usually granted on application, while Standard access review takes around 10 business days. Meta requires app creation and ad account permission grants. Audit your UTM naming convention: does the campaign ID in your UTMs match the campaign ID returned by the API? Mismatches make attribution

impossible. Also decide now where net revenue contribution comes from: MT4/MT5 spread data via Manager API, a manual monthly finance export, or a BI tool. The daily spend ingestion itself is built in Week 5, alongside the platform events, so that it is production-ready for Week 7.

0F. Expected VIP Tagger, deploy early as rules-only. The full tagger is a Week 9 build, but two earlier automations (Failed Deposit Rescue in Week 3, Early Churn Trigger in Week 8) and, in Week 10, the Losing-Streak Intervention and the VIP branch of the Withdrawal Trust Playbook reference the Expected VIP tag for escalation logic. Without it, they have no VIP signal to act on. Deploy a rules-only version now: `expected_vip = true` if FTD amount is at or above a threshold, typically the top 15 to 20% of historic deposit sizes; before a first deposit exists, the Failed Deposit Rescue tests the attempted amount against the same threshold. Set `expected_vip_source = 'rules_v1'` so you can distinguish rules-tagged from model-tagged clients later. One developer, half a day.

0G. Compliance sign-off plan. Book the compliance work the checklist in Appendix E requires, so that it does not become a Week 3 blocker: a lawful-basis record and legitimate interest assessment for the event stream, a data protection impact assessment for the rules-based VIP flag (live from Week 0) and for lead routing (Week 6), and financial-promotion review of every outbound template drafted in 0B before it is submitted to Meta. Half a day for the DPO or compliance lead to scope; the reviews themselves run in parallel with Weeks 1 to 5.

> **Week 0 exit gate.** Do not begin Week 1 until all seven checks are resolved or have a documented workaround. Unresolved items from Week 0 become Week 1 blockers.

Phase 1, weeks 1 to 4: foundation and immediate revenue recovery

Objective: install the event stream in stages, CRM first, then PSP and KYC, then MT4/MT5. Plug the two revenue leaks that recover spend already paid, failed deposits and KYC drop-off, before building routing or measurement intelligence. This phase should recover more revenue than it costs to build.

Week 1: CRM events and SLA infrastructure

Developer 3 to 4 days, CRM admin 2 days. Five working days, CRM integration only. Do not attempt PSP or MT4/MT5 this week.

Wire `lead_created` with source, campaign ID and geo captured and locked at creation, and implement a field lock so source cannot be overwritten by a rep after first write. Wire `lead_assigned` with timestamp and rep ID emitted the moment ownership is set; verify it fires on reassignment too. Wire `first_contact_attempt` and `first_contact_success` with channel, timestamp and rep, and build the CRM workflow that makes logging the path of least resistance rather than an extra step. If you run an IB channel, wire `ib_linked_to_lead` the same week and lock `ib_id_first_touch` with the same field rule. Set up the Event Stream Monitor for CRM events only. Build the SLA breach alert and a weekly SLA hit-rate report by rep. Until routing tiers exist in Week 6, apply a single interim one-hour SLA to every lead, and agree the tiered policy with the Head of Sales now so that Week 6 only switches it on: Tier A five minutes, Tier B one hour, Tier C same business day. Then manually trace 10 recent leads through the new event stream. Every gap is a bug. Fix before Week 2.

End-of-week measure: CRM event stream live and logging. SLA policy defined and communicated to the sales team. The 10-lead manual trace passes with no gaps in CRM-layer events.

The most common Week 1 discovery. The CRM overwrites `lead_source` when a rep edits the contact record. In one audit, 38% of funded accounts had their source overwritten to "Direct" because reps were cleaning up records manually. Six months of attribution data was corrupted. The fix is a single CRM field configuration: make source read-only after first write. It takes 30 minutes. Finding it takes knowing to look.

Week 2: PSP and KYC events

Developer 3 days, Payments/Ops 1 day for PSP mapping, IT half a day for the eKYC webhook. Four working days, plus half a day to a day per PSP beyond the first.

Wire `deposit_attempted` with amount, method, PSP and currency, one webhook integration per PSP. Wire `deposit_failed` with failure category using the code mapping table from Week 0; if the PSP does not distinguish abandoned checkout from

an active failure, treat a session timeout after `deposit_attempted` with no success within 30 minutes as abandoned. Wire `deposit_success` with amount, method, PSP and timestamp, since this event triggers the downstream retention automations. Wire KYC state transitions from your eKYC provider, building a reason classifier now if the provider returns unstructured text, because the Week 4 concierge depends on structured reasons. Extend the Event Stream Monitor to cover both.

End-of-week measure: deposit events flowing, `deposit_attempted`*,* `deposit_failed` *with category and* `deposit_success`*. KYC state transitions flowing from the eKYC provider. Monitor alerts configured for every new event type.*

Realistic timeline check. Three PSPs plus one eKYC provider is four separate webhook integrations in one week, each with its own authentication, payload schema and retry logic. At 0.75 days each that is three days of integration work before testing, error handling or monitoring. If your developer is also the CRM admin, this week runs long. Prioritise `deposit_failed` and `deposit_success` first: they unblock Week 3. KYC events can slip into Week 3 without blocking the critical path, but they must be stable before the Week 4 concierge is built.

Week 3: Failed Deposit Rescue

Payments/Ops 2 days, developer 1 day, Sales 30 minutes to review escalation rules. Three working days. The messaging channel must be approved before this week.

Build the bank decline playbook: trigger immediately, message one acknowledges the failure and lists two or three alternative payment methods with direct links; message two, two hours later if no retry, offers to connect with a payments specialist; human task if no action after 24 hours. Build the PSP error playbook: acknowledge, explain it is a technical issue and not their card, prompt retry, and auto-create a support ticket in parallel. Build the abandoned checkout playbook: trigger 30 minutes after `deposit_attempted` with no FTD and no failure event, message "your account is ready and waiting" with a direct link back, human call task after two hours. For VIP escalation, skip automated sequences entirely and create an immediate human task. Cap at two automated messages per failure event.

End-of-week measure: deposit failure rate as a share of attempts. Recovered deposit rate as a share of failures, targeting 40%+ within 30 days of going live.

Why this is Week 3, not Week 8. At the benchmark 200 FTDs per month and 1.2 attempts per FTD, about 40 deposit attempts fail. With an active rescue playbook, 40 to 60% are recoverable: 16 to 24 clients per month. At $800 NRC that is $12,800 to $19,200 per month recovered from acquisition spend already paid. Build this before lead scoring and before routing: it recovers existing spend rather than improving future spend, and it needs only the Week 2 events.

Week 4: KYC Completion Concierge

Ops/CS 2 days, developer 1 day. Two to three working days. Requires stable KYC events from Week 2.

Stall trigger one, no submission after `kyc_started`: fire two hours after with a document checklist specific to the client's jurisdiction; at 24 hours, "we are holding your account open" with a direct link back; at 48 hours, a human CS task. No further automated messages. Stall trigger two, needs-more-info or rejected: fire within 30 minutes, and use the provider's reason category to tell the client exactly what is missing. "Your ID document was unclear, please resubmit with better lighting" outperforms "your KYC requires attention" by a wide margin. If reason codes are unstructured, use the Week 2 reason classifier to map to one of four categories: document quality, wrong document type, ID mismatch, pending review. Stall trigger three, approved but not funded: at 24 hours after `kyc_approved` with no `deposit_attempted`, send the funding guide with the payment methods for the client's geo; at 72 hours, a sales task. Tag every stall with its reason category; this data feeds tier recalibration later.

End-of-week measure: KYC completion rate. Average time from `kyc_started` *to* `kyc_approved`*. Reason distribution for stalls. Phase 1 combined revenue impact estimate: $24,000 to $36,000 per month.*

Where message content matters more than timing. Unassisted KYC completion runs at 35 to 45% in the brokerages we audit. An active concierge with specific, reason-aware messaging typically lifts this by 8 to 12 percentage points. At the benchmark 700 KYC starts per month and a 10-point uplift, that is 70 additional completions; at the 25% funding rate for recovered completers, 17 to 18 additional FTDs per month, or about $14,000. The key variable is specificity. A generic "complete your KYC" recovers about 3 to 4% of stalls. A message that names the exact document missing and explains why recovers 10 to 15%.

Phase 1 KPIs to have established: event stream coverage of the CRM, PSP and KYC event types at 85%+ of records, with the Event Stream Monitor's missing-event rate heading toward its under-1% target; SLA hit rate, where even 40% is a valid baseline because you need a number and not an opinion; deposit failure rate and recovered deposit rate, targeting at least 40% of failures recovered; KYC completion rate before and after the concierge; and the Expected VIP tag rate from the Week 0 rules version. Combined revenue impact estimate: $24,000 to $36,000 per month.

Phase 1 cost at the benchmark: about nine developer days and eight days of CRM, payments and CS time, plus WhatsApp conversation fees that run to the low hundreds of dollars a month at 1,000 registrations. Against $24,000 to $36,000 a month recovered, the phase pays for itself inside its first month.

Phase 2, weeks 5 to 8: platform events, measurement and routing

Objective: complete the event stream with MT4/MT5 data, build routing intelligence, and stand up the measurement layer once you have six weeks of clean event data. The Profit Truth Agent moves to Week 7, not Week 5, because it needs six weeks of cohort data and production-ready ad spend integrations.

Week 5: MT4/MT5 events and the complete Event Stream Monitor

Developer 3 to 4 days, IT 1 day if API negotiation is required. Four working days with direct access, six or more if polling is required.

Wire `account_provisioned`, which starts the activation clock. Wire `first_login`, critical for the Week 8 Activation Coach. Wire `first_trade`, which closes the activation funnel and starts the retention clock. Wire the 3-trading-day inactivity signal as a scheduled daily job checking for clients who have placed a first trade and then gone three consecutive Monday-to-Friday trading days without one, excluding public holidays and broker non-trading days; this requires a trading calendar, so build or import one now. Complete the Event Stream Monitor across MT4/MT5 events. Build the daily ad-spend ingestion from the Google Ads and Meta Marketing APIs set up in Week 0, joined to the event stream by campaign ID, so that Week 7 has production data. Then manually follow five clients from click to lead to FTD to first trade. Every gap is still a bug at this stage.

End-of-week measure: provisioning, `first_login` *and* `first_trade` *all flowing. The 3-trading-day inactivity job running with correct calendar logic. Daily ad spend landing by campaign ID. The full event stream complete: click to lead to KYC to FTD to first trade, all traceable.*

> **The MT4/MT5 white-label problem.** In a typical white-label setup the provider controls the Manager API. The broker can see trades in the admin panel but cannot programmatically query account-level events. Getting access granted requires a written request, sometimes a contract amendment, and one to three weeks of back-and-forth. If you started that conversation in Week 0 you should have access by Week 5. If you did not, Week 5 becomes a waiting week. While you wait, build the daily report import. It is not as good as real-time webhooks but it unblocks Weeks 6 to 8.

Week 6: Lead Scoring and Routing V1, rules-based

Sales Ops/CRM admin 2 days, Head of Sales 1 day to define tier criteria. Three working days, no ML required.

Define A/B/C tiers using signals available at `lead_created`. Tier A is a high-intent source such as organic search or brand paid, plus a funded geo and a deposit indicator in the registration form. Tier B is a mid-intent source with a similar geo. Tier C is everything else. Route Tier A to your two best closers with a five-minute SLA, logging every routing decision with the signal values used. Route Tier B round-robin with a one-hour SLA, and Tier C to a same-business-day attempt, then a self-serve queue and nurture sequence. Switch on the tiered escalation ladder from the SLA Enforcer specification, replacing the interim one-hour SLA from Week 1. Build the weekly leakage map as the output of routing data; this is not a separate build, it is a SQL query against events you already have. If you run an IB channel, build the IB Lead Router alongside: route IB-sourced leads to the IB-specific desk with the SLA set by IB tier (Tier 1 IBs fastest), notify the IB relationship manager when Tier 1 or 2 IB leads arrive, and log `ib_id` on every decision.

End-of-week measure: FTD conversion rate by lead tier. Time to first contact by tier. Leakage at each funnel stage by tier. The weekly leakage map published.

Why rules first matters. Tier A leads, roughly 25 to 30% of volume, convert at 35 to 45% FTD when routed to the right rep with the right SLA. Tier C without routing converts below 5%. At the benchmark 1,000 leads per month, 250 Tier A leads at a 40% FTD rate is 100 FTDs from the top quarter of your pipeline. The data-informed re-tiering in Week 12 will improve on this, but the rules-based version delivers most of the lift immediately. Brief your two best closers on what a Tier A lead looks like and why they are receiving it: routing only compounds if the receiving rep treats it differently.

Week 7: Profit Truth Agent V1

Developer 3 days, CMO and CEO/COO 1 hour to define the memo format. Three working days, requiring Week 0 ad spend integrations to be production-ready.

Build CPA funded by campaign, geo and entity, joining ad spend to funded accounts via campaign ID. Every funded account must be traceable to a specific spend line; if any cannot be attributed, that is an event stream gap to fix. Build true C-ROAS to date: net revenue contribution divided by spend, by channel. Not registrations. Not deposits. NRC. Build mature-gated C-ROAS, where only sufficiently aged cohorts get a final figure; at Week 7 you have six weeks of data, so only the Week 1 and Week 2 lead cohorts are mature enough for a 30-day C-ROAS, and immature cohorts get forecast C-ROAS with confidence bands. Create the scale, hold or kill memo format: one page, three columns, distributed in the Monday leadership meeting from Week 8 onwards. Build the attribution coverage dashboard and the Attribution Integrity Monitor's alerts: click-ID capture rate, source overwrites, and an unknown-source alert at a 10% threshold.

End-of-week measure: CPA funded for at least three channels, not blended. True C-ROAS for the Week 1 and Week 2 cohorts at 30-day maturity. Attribution coverage 85%+ and unknown-source rate below 10%. The scale, hold or kill memo distributed.

Example Week 7 output. Channel A (Google Brand): CPA funded $280, 30-day C-ROAS 1.4x, **scale**. Channel B (Meta Lookalike): CPA funded $520, 30-day C-ROAS 0.6x, forecast 180-day 1.1x, **hold**, cohort immature. Channel C (Display Retargeting): CPA funded $890, 30-day C-ROAS 0.3x, forecast 180-day 0.6x even in the upside case, no improvement trend, **kill**. At Week 7 cohorts are only six weeks old, so 90- and 180-day figures are forecasts, not confirmed. A kill still needs a forecast: it is the upside case staying below break-even, not the 30-day number, that justifies it.

Week 8: Early Churn Trigger and Platform Activation Coach

Retention/CS 2 days, CRM admin 1 day. Three working days, requiring stable MT4/MT5 events from Week 5.

Platform Activation Coach: trigger on `account_provisioned` with no first login after 48 hours, and on first login with no first trade after 72 hours. Keep it to two automated touches maximum; the third is a human CS task. Content is practical, not promotional. Early Churn Trigger: fire at three consecutive trading days with no trade after the first trade, using the Week 5 calendar logic. The first touch is light and non-pushy; assume the client is still setting up. Cap total automated post-FTD touches at four per client in the first 14 days, because automated over-communication in that window causes more churn than no communication. Escalate Expected VIPs to a VIP desk task instead of an automated message.

End-of-week measure: post-FTD 7-day active rate, the share of FTDs who place at least one trade in their first seven days. Median time to first login and to first trade. Reactivation rate from the 3-day inactivity trigger.

The activation numbers. At the benchmark 200 FTDs per month, 35 to 50% will not place a trade within seven days of funding: 70 to 100 clients per month in the risk window. A well-executed two-touch activation sequence converts 15 to 25% of those to a first trade: 10 to 25 additional active traders per month. These clients have already paid the CPA and already funded, so the cost of the touch is near zero compared to the NRC recovered. Test the 3-trading-day trigger with five synthetic clients in different timezones before going live: calendar logic bugs are the most common cause of it firing late or incorrectly.

Phase 2 KPIs: full event stream traceable from click to first trade with clean IDs; CPA funded by campaign, geo and entity shared weekly; true C-ROAS for all channels with cohorts aged 30+ days, with the scale/hold/kill memo in the Monday leadership meeting; FTD conversion by lead tier showing weekly tier differentiation; attribution coverage 85%+ with unknown-source below 10%; and a post-FTD 7-day active rate baseline with first improvement from activation sequences.

Phase 2 cost at the benchmark: about seven developer days and seven days of IT, sales operations and retention time, plus the event store, typically $100 to $500 a month on a managed database or warehouse at this volume. Activation alone returns $8,000 to $20,000 a month.

Phase 3, weeks 9 to 12: intelligence and compounding

Objective: upgrade the rules-based work from Phases 1 and 2 with prediction and behavioural signals. Add the IB economics layer and VIP retention infrastructure. By Week 12 the system compounds, and the 90-day review tells you exactly where to invest in the next quarter.

Week 9: Expected VIP Tagger, full model upgrade

Developer 2 days, CEO/COO and retention lead 1 hour for threshold calibration. Two to three working days.

Upgrade from the Week 0 FTD-amount-only rule to a multi-signal model: FTD band weighted highest, plus time to first login (faster equals higher intent), geo and source channel quality. Assign a confidence score from 0 to 1 and set `expected_vip = true` above the threshold, so that the downstream automations keep reading the same boolean. Recalibrate the threshold using eight weeks of FTD data and four weeks of first-login data: which FTD and behaviour combinations are tracking toward Proven VIP (Appendix C) on the 30-day criteria? Target 60 to 70% expected-to-proven conversion once cohorts reach 90 days; below 40% means thresholds are too loose, above 85% too tight. All downstream automations already check the flag from Week 0, so the upgrade improves flag quality with no rewiring. Add `expected_vip_source = 'model_v1'`.

End-of-week measure: Expected VIP tag rate. Expected-to-proven conversion at 30 and 60 days on the proxy criteria. Confidence score distribution, where scores clustering at 0.5 are a sign of a weak model.

Why the upgrade matters. The Week 0 rules-only version tags on FTD amount alone, which typically over-tags by 30 to 40%: high-deposit clients who are one-time depositors and never trade. Adding time to first login is a strong predictor: a client who logs in within two hours of funding is three to four times more likely to become an active trader than one who logs in after 48 hours with the same deposit size. At the benchmark 200 FTDs per month, a 15 to 20% tag rate is 30 to 40 Expected VIPs a month, and the difference between a 60% and a 35% expected-to-proven rate is eight to ten misdirected clients: roughly 40 to 50 VIP desk hours per month redirected from low-probability to high-probability clients.

Week 10: VIP Losing-Streak Intervention and Withdrawal Trust Playbook

VIP Desk/Retention 2 days, Ops 1 day, developer 1 day for the equity data pipeline. Four working days.

Critical path. The Losing-Streak Intervention requires daily P&L and account equity history per client, from MT4/MT5 via Manager API or the daily report import built in Week 5. If access is still restricted and no import exists, build the Withdrawal Trust Playbook only, which needs no equity data, and defer the intervention to a future sprint.

Losing-Streak Intervention: monitor Proven and Expected VIP clients only for drawdown. Threshold: three consecutive losing days **and** drawdown beyond 20% of account equity at day start. Action is a VIP desk task, human outreach only, scripted with compliance and recorded. Build the equity data pipeline as a daily job storing account equity and daily P&L per VIP client; MT4 keeps no per-account daily equity history and MT5's end-of-day records are only reachable through Manager API access you may not have, so budget half a day to a day for the storage layer. Withdrawal Trust Playbook: trigger on `withdrawal_requested` for any client; within one hour send an acknowledgement with a realistic, specific timeline ("2 to 3 business days", not "soon"); at 24 hours still pending, a status update; at 48 hours, an Ops escalation with a retention notification; at completion a thank-you; for clients who withdrew less than 70% of equity, an optional re-engagement message seven days later; for VIP clients, a personal call from the retention desk. Log withdrawal reasons, because patterns in them are product or service problems worth fixing.

End-of-week measure: withdrawal acknowledgement time, target under 60 minutes. VIP churn rate, the share of VIP clients withdrawing all of their equity within 90 days of FTD. Post-withdrawal re-deposit rate at 30 days.

The easiest win in this section. Most brokerages send no communication when a withdrawal is requested unless the client chases. A single automated acknowledgement within 60 minutes, confirming receipt and giving a specific timeline, reduces complaint escalations measurably and improves re-deposit rates. It costs one afternoon to build.

Week 11: IB Partner Retention Assist and IB Economics Engine

IB Manager and Finance 2 days, developer 1 day. Three working days. Skip entirely if you have no IB channel.

IB Economics Engine: weekly profitability by IB and tier. Gross revenue from IB-sourced clients minus commissions paid equals net IB contribution, and this number drives tier reviews. Add a commission accuracy check that verifies `commission_accrued` events match agreement terms, flagging any discrepancy above 2% for manual review. Define the net contribution and cohort quality thresholds that trigger an upgrade, a restructure or an offboard conversation, on a quarterly cadence. IB Partner Retention Assist: trigger on elevated inactivity or withdrawal rates in an IB's cohort, for example a 7-day active rate 30% below that IB's historical average, and create a task for the IB Manager to contact the partner before the problem compounds. Complete the IB Attribution Monitor the same week: a daily audit that alerts on any client mapped to two IBs and any IB-linked lead that has lost its `ib_id`, quarantining the affected cohort before the commission run.

End-of-week measure: net IB contribution by partner available and shared with the IB Manager weekly. Commission accuracy rate. IB cohort FTD rate against the non-IB benchmark. The IB attribution audit running daily.

The most common IB Economics discovery. When the engine runs for the first time, the top IB by volume is typically third or fourth most profitable by net contribution after commission. A smaller IB with a better-quality cohort and a lower commission rate delivers 40% more net NRC per referred FTD. Without this, the IB Manager optimises for volume because that is all that is visible. With it, the tier review conversation with the high-volume, low-quality IB has a number behind it, not a feeling.

Week 12: Lead Scoring upgrade and the 90-day review

Developer 1.5 days, CMO, Head of Sales and CEO/COO for a two-hour review session.

Upgrade Lead Scoring from rules-based to data-informed using six weeks of logged routing decisions and ten weeks of lead and FTD events. Which source, geo and signal combinations actually converted at Tier A rates? Which ones you assumed were Tier A were actually Tier B? Update the tiers accordingly. Run the 90-day scoreboard review with the full leadership team. Every metric must have a number; if any cell is still an estimate, you have an event stream gap to add to the Q2 backlog. Identify the single highest-impact improvement the data reveals, and make it the Q2 priority. Finally, evaluate ML readiness for the three AI agents: moving from rules to ML typically requires four to six months of clean labelled data at meaningful volume, so if you are below 200 labelled conversions, stay rules-based for now.

End-of-week measure: the day-90 scoreboard below, with a number in every cell.

Phase 3 cost at the benchmark: about six developer days and five days of retention, operations, IB and finance time, and no new infrastructure beyond the equity pipeline's storage.

What 10 weeks of data typically reveals. One channel that looked expensive by CPA has a 2x better NRC per client than any other, so scale it. One geo converts well at KYC but has poor post-FTD retention, which is a product-market fit issue rather than a funnel issue. One rep whose routing-assigned leads convert at 2x the desk average, which is a coaching and pairing opportunity. One KYC stall reason that accounts for 40% of all stalls, which is a document guidance fix that takes one afternoon. None of these are visible without the event stream.

The day-90 TRACE scoreboard

Every cell below must have a number at the Week 12 review. If any cell is still an estimate or a narrative, you have an event stream gap.

METRIC	WHAT GOOD LOOKS LIKE
Attribution coverage	85%+ of leads with a valid source and campaign persisted to FTD. Unknown-source rate below 10%.
CPA funded	Segmented by campaign, geo and entity. Not a blended average.
True cohort ROAS	At least three channels with 30-day-aged cohorts. One clear scale signal, one clear kill signal.
SLA hit rate, Tier A leads	70%+. Below 50% means the policy is not enforced: a management issue, not a tech issue.
Deposit failure rate	Baseline established from Week 2 data.
Recovered deposit rate	40%+ of failed deposits rescued within 30 days of go-live.
KYC completion rate	8 to 12 percentage points above the pre-launch baseline.
Post-FTD 7-day active rate	Measurable improvement from the pre-Week-8 baseline. Any improvement is retained NRC.
3-day inactivity reactivation	A measured share of triggered clients returning to trade within 7 days of the intervention; 15 to 25% is the range we see in audits.
FTD rate by lead tier	Tier A converting at 3x+ the rate of Tier C. If not, the tiers are not differentiated.
Net IB contribution	By partner, after commission. At least one tier review decision made.
Expected VIP accuracy	Expected-to-proven conversion tracking toward 60 to 70% on the 30-day proxy criteria; the 90-day figure lands between day 120 and day 180. Below 40% means retune the thresholds.
Withdrawal acknowledgement	Under 60 minutes from request to client confirmation.
VIP churn rate	Declining trend from the Week 9 baseline.

The honest question at day 90: how much of this revenue was recoverable before the system existed? At the benchmark 1,000 leads per month, the combined estimate across the Phase 1 and Phase 2 automations alone is $32,000 to $56,000 per month in recovered or protected revenue: $12,800 to $19,200 from deposit rescue, $11,200 to $16,800 from the KYC concierge and $8,000 to $20,000 from activation, before counting speed-to-lead and routing gains. That number scales linearly with volume. It does not require a CRM migration. It requires clean events, week-zero pre-work completed properly, and the discipline to deploy in the right order.

If you try to deploy everything simultaneously, you will recreate the Operator Trap. TRACE works because it installs the compounding system in the correct order.

APPENDIX A

Minimum viable event stream checklist

The baseline definition of Truth in TRACE. If you cannot reliably capture these events with timestamps and stable IDs, you cannot measure leakage, tie spend to net revenue contribution, or deploy serious agents.

The rule: every event must carry a stable lead ID or client ID, a UTC timestamp, geo and entity, and source or campaign identifiers where relevant. If IB-sourced, every relevant event must also carry `ib_id`, `sub_ib_id` and tier fields.

1. Identity and attribution (click to lead). Ad click captured, with click IDs persisted to the session or user record · lead created, with source and campaign persisted · source and campaign persistence verified, meaning source does not get overwritten after registration, KYC or deposit.

2. IB attribution (IB channel only). IB linked to lead, with `ib_id` / `sub_ib_id` attached and immutable · IB attribution persistence verified · IB agreement and tier recorded, with version available for audit.

3. Sales execution. Lead assigned, with rep or desk ownership set and the SLA clock started · first contact attempt logged with channel, timestamp and rep · first contact success logged, meaning connected or replied rather than merely attempted · qualification call completed with outcome logged.

4. Onboarding. KYC started · KYC submitted · KYC decision with reason category · eKYC provider status synced into your event stream.

5. Funding. Deposit attempted, with method, amount, currency and PSP · deposit failed, with failure category, reason code and retry count · deposit successful, with amount, method and time.

6. Platform activation. Trading account provisioned, with credentials issued, server assigned and account type · first platform login · first trade, with asset class or symbol and optional leverage band.

7. Post-FTD retention signals. Three consecutive trading days with no trade after the first trade, Monday to Friday only · retention or reactivation touchpoint logged, with channel and intent.

8. Trust and lifecycle. Withdrawal requested, with amount and reason if available · withdrawal completed or failed, with time to complete and failure reason.

9. IB economics. IB commission accrued, with rule version, basis and amount · IB commission paid, with timestamp · IB dispute or open case logged.

Truth quality checks (strongly recommended). ID integrity, meaning consistent lead-to-client mapping with duplicates and merges tracked · attribution coverage tracked weekly · event completeness, with missing-event alerts such as an FTD with no deposit attempt logged · latency, with ingestion delay measured · audit trail, with event logs immutable or versioned, especially for regulated entities.

APPENDIX B

Broker attribution stack checklist

If your answer is no to any of these, you are optimising blind.

- We can connect spend to campaign to lead to funded customer.
- Campaign and source identifiers persist through registration, KYC and deposit without being overwritten.
- We have a reliable lead ID to client ID mapping, with duplicates and merges tracked.
- Deposit attempted, failed and successful events exist with reason codes from the PSP.
- MT4/MT5 provisioning, login and first trade signals exist.
- We can calculate CPA funded by campaign, geo and entity.
- We can compute true C-ROAS to date and mature-gated C-ROAS at 7, 30, 90 and 180 days.

- We produce forecast C-ROAS at 90 and 180 days with confidence bands for immature cohorts.
- We do not mix IB network referrals into trader acquisition ROAS.
- We maintain an attribution integrity alerting system.

APPENDIX C

VIP definitions: expected versus proven

Expected VIP. A client who has funded and exhibits early signals suggesting above-average lifetime value, but whose full value has not yet been confirmed by trading behaviour over a sufficient period. Rules-based V1 (Week 0): FTD amount in the top 15 to 20% of historic deposit sizes for the entity. Scored V2 (the Week 9 model) adds time to first login; geo tier A (historically a higher lifetime-value cohort); source quality score above threshold; and early platform engagement (multiple instrument views within the first session and session length above median).

Proven VIP. A client who has demonstrated high value through actual trading behaviour over a sufficient measurement period, typically 30 to 90 days post-FTD. Example criteria, to be calibrated to your own cohort data: trading volume above X over 30 days; net revenue contribution above Y over 90 days; cumulative deposits above Z; trading frequency meeting a minimum threshold, such as at least three trading days per week.

> The Expected VIP tag enables premium retention treatment from day one, before the client has proven their value. The transition from Expected to Proven is your most important model validation signal. Track it monthly and recalibrate thresholds quarterly.

VIP desk treatment protocol. Expected VIPs are assigned to a dedicated retention rep, receive a priority manual touch within 24 hours of tagging, and are excluded from mass automated sequences. Proven VIPs get a dedicated relationship manager, personal outreach for major market events, priority processing for all requests, and a two-hour losing-streak intervention SLA.

APPENDIX D

Vanity ROAS versus true cohort ROAS: a worked example

This demonstrates the difference between vanity ROAS, which most brokers calculate, and true C-ROAS on net revenue contribution, which TRACE requires.

Scenario (hypothetical, one campaign-month). Campaign: Google Search. Geo: UAE. Ad spend $10,000. Registrations 200. FTDs 20, so a 10% FTD rate: this single campaign converts at half the benchmark funnel's rate, which is part of the point. CPA funded $500. Average FTD amount $1,500. Total deposit volume $30,000.

Vanity ROAS (what most brokers report). Deposit volume divided by ad spend: $30,000 / $10,000 = **3x**. This looks acceptable. But it is not ROAS. It is a volume ratio, and it says nothing about profitability.

True C-ROAS on net revenue contribution (the TRACE method).

Gross revenue from this cohort at 90 days (spread, commissions, swap fees)	$8,000
Rebates, bonuses and payment costs attributable to the cohort	$1,200
Net revenue contribution	**$6,800**
True C-ROAS = NRC / spend	**0.68x**, below break-even at 90 days

Forecast C-ROAS (for immature cohorts). Historical value curves for similar geo and source cohorts show a 90-day C-ROAS averaging about 0.7x but a 180-day C-ROAS averaging 1.4x. Forecast for this cohort at 180 days: base 1.4x, downside 1.0x, upside 1.8x. **Decision: hold, not kill.** The channel is below break-even at 90 days but forecast to be profitable at 180. Flag for re-evaluation at day 120.

> A channel that looks profitable at vanity ROAS may be destroying value at true C-ROAS. A channel that looks unprofitable at 90-day true C-ROAS may be highly profitable at 180 days. Never make scale or kill decisions on immature cohorts without forecast C-ROAS.

APPENDIX E

Compliance readiness checklist

This checklist is for compliance officers, DPOs and legal teams reviewing a TRACE deployment. It is not a legal opinion and does not substitute for qualified regulatory counsel. It is designed to give compliance a structured starting point, so that their first response to a deployment proposal is a set of specific, answerable questions rather than a blanket hold.

TRACE is designed to be compliance-legible. Every automation fires on a documented event trigger, produces a logged output, and can be paused or overridden by a human. **No AI agent in TRACE makes a decision a human cannot see and override within the SLA.** They produce scores and signals; the automated actions they drive, routing a lead to a desk or a client to a VIP queue, are logged and reversible by the owner.

#	QUESTION	IF NO: IMPLICATION
1	Has a lawful basis been documented for each processing activity under the applicable privacy law (GDPR Article 6, UK GDPR, PDPL, POPIA)? Consent, legitimate interest and contractual necessity are distinct bases with different obligations.	Processing is unlawful. Automated communications and profiling are high-risk without a documented basis.
2	Where the basis is legitimate interest, has a Legitimate Interest Assessment been completed?	Legitimate interest claims are regularly overturned by regulators without a documented LIA. High risk for EU and UK deployments.
3	Where consent is required for outbound marketing, is it freely given, specific, informed and unambiguous, with a withdrawal mechanism in place?	Outbound automated messages without valid consent violate GDPR Articles 6 and 7, the ePrivacy Directive or PECR, and their equivalents. Templates must be reviewed.
4	For the three AI agents, do any outputs constitute automated individual decision-making with legal or similarly significant effect under GDPR Article 22?	If yes, the right to human review, explanation and challenge must be implemented. This applies to routing that materially affects access to financial services.
5	Have Data Processing Agreements been executed with every third-party processor, including CRM, eKYC, PSPs, messaging platforms and analytics tools?	DPAs are a legal requirement under GDPR Article 28. Missing DPAs create direct regulatory exposure.

#	QUESTION	IF NO: IMPLICATION
6	For high-risk profiling, particularly the Expected VIP Tagger and Lead Scoring Agent, has a Data Protection Impact Assessment been completed?	DPIAs are mandatory under GDPR Article 35 for systematic profiling. Deploying without one creates direct enforcement risk.
7	Have all outbound templates been reviewed against applicable financial promotion rules? FCA COBS 4, MiFID II standards, SCA, DFSA and FSRA approval requirements in the UAE.	Automated messages that constitute financial promotions require pre-approval in most regulated jurisdictions. Template review is non-negotiable.
8	Does the VIP Losing-Streak Intervention, which monitors drawdown and triggers human outreach during losses, constitute a regulated activity such as investment advice?	Proactive contact during a losing streak may cross into regulated advice in some jurisdictions. The human outreach script must be reviewed.
9	Is an immutable audit trail maintained for all automation-triggered communications, routing decisions and scoring outputs?	Without an audit trail, outputs cannot be demonstrated to regulators. Immutable logs are a minimum requirement for any regulated deployment.
10	For cross-border data flows, is a valid transfer mechanism in place: Standard Contractual Clauses, an adequacy decision, or equivalent?	Cross-border transfers without a valid mechanism violate GDPR Chapter V. Cloud infrastructure jurisdiction must be audited.
11	For UK and EU entities, has the Churn Engineer layer been reviewed against the FCA Consumer Duty and the equivalent CySEC conduct expectations? In particular: is any retention contact timed to a client's losses or to an intended withdrawal, and could it be read as discouraging the client from acting in their own interest?	Retention outreach that leans on a losing client, or delays a withdrawal, is the conduct the Duty's consumer-support and consumer-understanding outcomes target. The script, the timing and the human reviewer must be documented.
12	Are the phone, WhatsApp and chat channels the Route and Activate layers run on recorded and archived in line with MiFID II Article 16(7) and FCA SYSC 10A, with the record linked to the client ID?	Off-channel messaging on personal devices is an active enforcement theme. Every automated or human client communication that could relate to a transaction must be captured.
13	Does the registration flow collect an explicit, recorded request for contact, so that the five-minute call is a solicited communication rather than a cold call under FCA COBS 4.8 and its equivalents?	Without the recorded request, speed to lead becomes unsolicited real-time promotion of a CFD, which is restricted. The consent wording and its timestamp belong in the event stream.

Recommended sign-off process. Complete this checklist before any automation goes live. For each "no" or "partial", assign an owner and a resolution date. Legal counsel reviews all outbound message templates before activation, and the retention and losing-

streak scripts before Phase 3. DPIAs completed before the rules-based Expected VIP flag goes live in Week 0 and before lead routing goes live in Week 6 (see 0G). Audit trail architecture confirmed before Week 1 of the rollout.

Jurisdiction reference

High-level reference only. Regulatory requirements change; always verify current obligations with qualified counsel in each jurisdiction.

JURISDICTION	RISK	PRIMARY CHECKLIST ITEMS
CySEC / EU	Amber	1–6 (GDPR lawful basis, LIA, consent, Article 22, DPA, DPIA); 7 (MiFID II financial promotions); 9–10; 11–13 (conduct expectations on retention contact, MiFID II Article 16(7) recording, solicited contact).
FCA / UK	Amber	1–6 (UK GDPR); 7 (COBS 4, all automated templates require review); 8 (advice boundary); 9–10; 11–13 (Consumer Duty, SYSC 10A recording, COBS 4.8 cold calls).
UAE (SCA / DFSA / FSRA)	Amber to red	1–3 (PDPL consent onshore; the DIFC and ADGM data protection laws in the free zones); 7 (SCA promotion and introduction rules; DFSA and FSRA financial promotion rules); 8 (advice boundary particularly sensitive); 9–10.
FSCA / South Africa	Amber	1–3 (POPIA, prior consent for marketing); 7 (advertising rules under the FAIS General Code of Conduct, section 14); 8 (advice boundary under FAIS); 9–10.
ASIC / Australia	Amber	1–3 (Privacy Act; Spam Act consent for electronic marketing); 7 (design and distribution obligations; the CFD product intervention order); 8 (general advice warning); 9–10.
Offshore (SC / VU / VC / MU)	Amber	The local conduct burden is light, but the exposure follows the client: apply 1–3 conservatively and 11–13 for EU and UK client segments regardless of broker licence, because reverse-solicitation and financial-promotion rules follow the resident. 9 (audit trail) and 10 (data transfers) always apply. St Vincent and the Grenadines has required evidence of licensing elsewhere since 2023.

TRACE is an operational framework, not a compliance waiver. This checklist is a starting point for compliance engagement, not a substitute for legal review. Jupiter Tech recommends completing it with your DPO and external legal counsel before deploying any component in a regulated jurisdiction.

FAQ: the TRACE framework for CFD brokers

What is the TRACE framework?

TRACE is a revenue engineering framework for retail FX and CFD brokerages, built by Jupiter Tech from operational audits of more than 30 of them. It treats the funnel from ad click to funded, trading client as an event stream, and deploys 16 components, three AI agents and 13 deterministic automations, in a fixed order so that each stage's data makes the next one work.

What does TRACE stand for?

Truth, Route, Activate, Churn Engineer, Economics. Truth is owning the event stream and attribution. Route is allocating sales attention by lead score and SLA. Activate is removing friction at KYC and the first deposit. Churn Engineer is intervening in predictable post-FTD risk windows. Economics is running the brokerage on net revenue contribution by cohort rather than on registrations or deposit volume.

How long does TRACE take to implement?

Twelve build weeks after a mandatory Week 0 of prerequisites, about 90 days of build time, with the first revenue recovery, failed deposit rescue, in Week 3. Elapsed time is usually 13 to 16 weeks, because Week 0's vendor approvals (MT4/MT5 Manager API access, WhatsApp business verification and ad-platform API access) take one to three weeks to come back.

Do I need a new CRM to run TRACE?

No. TRACE needs ownership of your event stream, not a new system of record. If your CRM exposes webhooks and endpoints, stream the events out and push tasks back in. If it does not, build a broker-owned event store fed by the CRM, the PSPs, the KYC provider, the trading platform and the ad platforms. Migrate only if the vendor blocks both paths.

Which TRACE components use AI, and which are automations?

Three use machine learning: the Profit Truth Agent (forecast cohort ROAS), the Lead Scoring and Routing Agent (upgraded from rules to a model once it has four to six months of labelled data) and the Expected VIP Tagger (a behavioural scoring model). The other 13 are deterministic automations: each fires on an event, applies a rule and produces a logged action. All three AI agents ship first as rules.

What is speed to lead, and why does it matter for a broker?

Speed to lead is the time between a prospect registering and the first human contact attempt. InsideSales' 2021 study of 5.7 million inbound leads found conversion 8x higher when the first attempt came within five minutes rather than later the same day, yet 57.1% of first call attempts happened after more than a week. In the benchmark brokerage, slow contact costs 30 to 50 first-time deposits a month.

What is true cohort ROAS (C-ROAS)?

True cohort ROAS is the net revenue contribution of a cohort of funded clients divided by the acquisition spend that produced them, reported as final only once the cohort is old enough (mature-gated) and as a forecast with confidence bands before then. Vanity ROAS divides deposit volume by spend and says nothing about profitability: a campaign can show 3x vanity ROAS and 0.7x true C-ROAS at 90 days.

How much revenue does funnel leakage cost a CFD broker?

On the paper's benchmark model, a brokerage with 1,000 registrations a month, 200 first-time deposits and $800 of net revenue contribution per funded client over 180 days loses roughly $600,000 to $900,000 a year to three fixable leaks: slow first contact ($288,000 to $480,000), KYC drop-off ($134,000 to $202,000) and failed deposits ($192,000). The figures scale linearly with volume, and every calculation is shown.

What is an Expected VIP?

An Expected VIP is a newly funded client whose early signals, first deposit size, time to first login, geo and source quality, suggest above-average lifetime value before trading history can confirm it. The tag routes the client to premium retention treatment from day one. A Proven VIP has confirmed that value through 30 to 90 days of trading behaviour, and the expected-to-proven conversion rate is the model's main accuracy check.

Is TRACE compatible with GDPR and financial promotion rules?

It is designed to be: every automation fires on a documented trigger, logs its output and can be paused or overridden by a human, and no AI agent makes a decision a human cannot see and reverse. Compliance still has to clear it. Appendix E is a thirteen-question readiness checklist covering lawful basis, legitimate interest assessments, consent, Article 22, processor agreements, DPIAs, financial promotion review, the advice boundary, audit trails, cross-border transfers, Consumer Duty, communications recording and solicited contact, with a jurisdiction table for CySEC, the FCA, the UAE, the FSCA, ASIC and offshore licences.

Sources

- **InsideSales / XANT, Lead Response Study 2021.** 55 million sales activities across 5.7 million inbound leads at more than 400 companies. Conversion rates are 8x higher when the first attempt is made within five minutes rather than between five minutes and 24 hours; 57.1% of first call attempts occur after more than a week. Cited in Section 2 and the leakage calculation.
- **Fenergo, KYC and Onboarding Trends 2024.** Survey of more than 450 C-suite executives at investment and commercial banks; 67% reported losing clients to slow or inefficient onboarding. Cited in the leakage calculation.
- **Signicat, The Battle to Onboard 2022.** Survey of 7,600 consumers across 14 European countries; 68% had abandoned a financial application during onboarding, with time to apply, the amount of personal information required and changing their mind the main reasons. Cited in the leakage calculation and Automation 6.

- **Jumio, How to Reduce Customer Onboarding Abandonment.** KYC/AML friction can drive onboarding abandonment as high as 70 to 80% in heavily regulated flows; a vendor estimate without a cited sample. Cited in the leakage calculation.
- **Checkout.com and Oxford Economics, Research on payment performance across the US, UK, France and Germany, September 2023 (2022 data).** Survey of 1,500 businesses and more than 8,000 consumers. 45% of consumers would not retry after a single false decline and 42% would never return; US businesses lost 2.1% of revenue to false declines in 2022. Cited in the leakage calculation.
- **ESMA, Product intervention measures on CFDs, 27 March 2018.** Leverage limits, margin close-out, negative balance protection, a restriction on incentives and a standardised risk warning; 74 to 89% of retail CFD accounts were found to lose money. Cited in Section 1.4.
- **FCA, Permanent restrictions on the sale of CFDs and CFD-like options to retail consumers, 1 July 2019.** Firms must stop offering monetary and non-monetary inducements to encourage trading; rules apply from 1 August 2019. Cited in Section 1.4.
- **Meta, WhatsApp message template review and Business Verification.** Templates must be approved before use and review can take up to 24 hours; business verification gates the higher messaging limits. Cited in Week 0.
- **Google Ads Help, Guidelines for importing offline conversions.** Offline conversions uploaded more than 90 days after the last click are not imported; 63 days for enhanced conversions for leads. Cited in Section 2.5.
- **Google Ads Help, How to adjust your conversions.** Conversion adjustments can be made within 54 days of the original conversion, and within 7 days for Smart Bidding to use them. Cited in Section 2.5.
- **Google Ads API, Access levels and permissible use.** Basic access requires Google Cloud brand verification and is often granted on application; Standard access review takes around 10 business days. Cited in Week 0.
- **Google Ads Help, About offline conversion imports.** Technical documentation for uploading offline conversion events to Google Ads. Cited in Section 2.5.
- **Google Ads Help, About enhanced conversions for leads.** Technical documentation for enhanced conversion matching using hashed first-party lead data. Cited in Section 2.5.

- **Meta Business Help Centre, About Conversions API.** Technical documentation for server-side event transmission to Meta. Cited in Section 2.5.

Compliance and regulatory references

- **European Data Protection Board, Guidelines 2/2023 on the technical scope of Article 5(3) of the ePrivacy Directive.** Which tracking techniques fall under the ePrivacy Directive's consent requirement. Cited in Appendix E.
- **legislation.gov.uk, UK GDPR Article 22: automated individual decision-making, including profiling.** The retained text of Regulation (EU) 2016/679 Article 22. Cited in Appendix E.
- **European Commission, Data protection obligations for business and organisations, including when a DPIA is required.** When a data protection impact assessment is mandatory under GDPR Article 35, including systematic and extensive profiling. Cited in Appendix E.
- **ESMA, MiFID II Article 24: general principles and information to clients.** The rule governing marketing communications and client information standards. Cited in Appendix E.
- **ESMA, ESMA reports on the application of MiFID II marketing requirements.** Supervisory convergence report on how national competent authorities apply the marketing rules. Cited in Appendix E.
- **FCA, Contract for differences.** FCA guidance for CFD providers on marketing, client categorisation and leverage. Background to Section 1.4 and Appendix E.
- **FCA Handbook, COBS 4.2: fair, clear and not misleading communications.** The standard for all client communications, including automated outbound messages. Cited in Appendix E.
- **FCA Handbook, SYSC 10A: recording telephone conversations and electronic communications.** The duty to record and retain communications that relate to transactions. Cited in Appendix E, question 12.
- **FCA Handbook, COBS 4.8: cold calls and other promotions that are not in writing.** When a real-time financial promotion may be made without an express request. Cited in Appendix E, question 13.

- **ESMA, MiFID II Article 16: organisational requirements.** Article 16(7), the recording of telephone conversations and electronic communications. Cited in Appendix E, question 12.
- **FCA, Consumer Duty.** The FCA's hub for the Consumer Duty, including fair value, consumer understanding and consumer support. Cited in Appendix E, question 11.
- **FCA, Outsourcing and operational resilience.** Obligations for regulated firms using third-party technology. Background to Appendix E, question 5.
- **ICO, Direct marketing and privacy and electronic communications.** Direct marketing rules under UK GDPR and PECR. Cited in Appendix E.
- **ICO, What does the UK GDPR say about automated decision-making?** UK GDPR Article 22 obligations for automated decision-making and profiling. Cited in Appendix E.
- **UAE Government, Data protection laws.** Official overview of the UAE Personal Data Protection Law and its scope. Cited in Appendix E.
- **Information Regulator (South Africa), Guidance notes.** Official POPIA guidance, including direct marketing obligations. Cited in Appendix E.
- **popia.co.za, POPIA section 69: direct marketing by means of unsolicited electronic communications.** The consent requirement for electronic direct marketing in South Africa. Cited in Appendix E.
- **South African Government, Government Gazette No. 25299, General Code of Conduct for Authorised Financial Services Providers and Representatives.** The FAIS General Code of Conduct, including the advertising requirements in section 14. Cited in Appendix E.
- **FSCA and Prudential Authority, Joint Communication 2 of 2025: cloud computing and data offshoring.** Cloud and cross-border data obligations for South African financial institutions. Cited in Appendix E.
- **Data Protection Office, Mauritius, The Data Protection Act 2017.** Personal data processing for entities operating under Mauritian regulation. Cited in Appendix E.
- **Republic of Seychelles, Data Protection Act, 2023.** Personal data processing for FSA-regulated entities in Seychelles. Cited in Appendix E.
- **Republic of Vanuatu, Data Protection and Privacy Act No. 13 of 2024.** Data processing obligations for VFSC-regulated entities. Cited in Appendix E.

Charlton Jones

Co-founder, Jupiter Tech

Co-founder of Jupiter Tech and of Convertico, the performance marketing and web agency that serves more than 40 clients. Built Passify, which automated trading strategies on behalf of retail traders, brokers and prop firms, more than 70 in all, and turned an obsession with how brokerages actually function into Jupiter Tech.

LinkedIn · **About Charlton**