


AN AGENT STACK WHITE PAPER

# An engine is not a car

*What an agentic AML/KYC compliance program for private markets actually requires.*



## KEY TAKEAWAYS

# Anthropic shipped an engine. A GP cannot run a program on one.

WHAT TO TAKE AWAY	3 points
01	In May 2026, Anthropic published a KYC screening agent as part of its Claude for Financial Services release. It is real, useful work: it reads investor documents, checks them against a firm's stated policy, and assembles a packet for a compliance officer. Think of it as an engine — genuine horsepower, the hard part done well.
02	A compliance program is the whole vehicle. It includes the parts every program has needed for years, well before anyone said <i>agent</i> , and a second set of parts the AI-native era demands. Most of those parts are deterministic software, careful workflow, and human judgment. None of them arrive in a markdown file.
03	Juniper Square has spent more than twelve years building the car. More than 2,000 GPs run on the platform, 750,000 LP profiles sit inside it, and 425+ fund administration professionals run real compliance programs every day. The compliance agent now in development bolts the new engine into that car. It makes the program faster. It does not replace it.

# 87%

of institutional investors have declined or reconsidered a fund commitment over AML/KYC concerns.

# 96%

now say the strength of a firm's program matters when they decide whether to invest.

SOURCE —  
CSC GLOBAL AML/KYC REPORT,  
NOVEMBER 2025

## OPENING

# The work sounds simple



## IN ONE SENTENCE

*Gather. Check. Then accept, decline, or ask for more.*

The idea of a compliance program is almost embarrassingly simple to describe. You gather data and documents from an investor. You check what they sent against a rulebook. Then you do one of three things: accept them, decline them, or go back and ask for more. That is the whole shape of it, and you could explain it to someone over coffee in about a minute.


The reality is a different animal. Collecting data and documents sounds trivial until you try to do it for investors who live in dozens of countries, subscribe to funds domiciled in dozens of jurisdictions, and arrive through every channel from a wirehouse feeder to a hand-scanned PDF. The applicable rules change with the jurisdiction of the fund, the type of the fund, and the asset class. They sit on top of data privacy and security regimes that differ by country and that a GP cannot afford to get wrong.

The same investor who is routine in one structure triggers an entirely different document set in another. As Joshua Demers, who leads investor services at Juniper Square and has run compliance in private markets for eighteen years, puts it, onboarding comes down to data and documents — and the reason it is hard is the sheer variability of what comes in the door.

Then there is the part that almost no one mentions until they have lived through it. On top of all that jurisdictional and regulatory complexity, you have to give the investor an experience that feels considered rather than punishing. That is its own design problem, and it is the one that separates a program a GP is proud of from one their LPs quietly resent.

So when Anthropic shipped a compliance artifact in May, the right reaction was neither to dismiss it nor to panic. The right reaction was to look closely at what it is, give it real credit, and be precise about everything it leaves for you to build. That is what this paper does.

- *The shape is simple. The variability is the whole job.*



WHAT ANTHROPIC SHIPPED

# An engine, and what it does well

The KYC Screener converts raw model capability into useful motion on a discrete piece of the job — and it does that well. Strip it to its essence and you have two genuinely useful primitives.

---

## IN PLAIN TERMS

 **FOUR WORDS TO  
PIN DOWN****Agent**

Software that uses an AI model to carry out a task on your behalf.

**Skill**

A detailed, reusable set of plain-language instructions for one specific job.

**MCP**

A standard way to plug an agent into an outside data source without custom engineering.

**Probabilistic vs. deterministic**

AI models can give slightly different answers to the same question; traditional code returns the *identical* answer every time. Hold onto this one.

Anthropic released the KYC Screener as part of Claude for Financial Services, a package of agent templates, skill bundles, and partner data integrations available to anyone on a paid Claude plan. Before going further, a few terms are worth pinning down, because they carry the whole discussion.

With those in hand, here is what the KYC Screener actually contains. It is built from three skills. The first reads a submitted investor package and pulls the relevant information into structured fields. The second takes the firm's AML/KYC policy, which the firm supplies as a plain-text file, and asks the model to judge the investor against it, calling out to a screening provider if the firm has wired one up. The third assembles a spreadsheet — an escalation packet a compliance officer can open and review. A capable team can have this running against sample documents in an afternoon, and the marginal cost is zero for firms already paying for Claude.

Four things about this release deserve to be said plainly, because they are all true and all to Anthropic's credit.

## TO ITS CREDIT

**It lowers the barrier to trying AI on real compliance work.** For a chief compliance officer who has watched AI from a distance, the KYC Screener is a concrete place to start. It is working code, not a pitch deck.

**It settles a question that was open a year ago.** Plenty of compliance professionals were still asking whether AI was serious about AML and KYC at all. Anthropic answering yes, and shipping something people can run, ends that debate.

**It rides improvements that lift everyone.** Better base models mean better document extraction, entity classification, and adverse-media summarization. Juniper Square's agents run on the same underlying models, so progress at the foundation helps the whole field at once.

**And Anthropic is honest about the limits.** The agent cannot reach a final determination on its own. It is built for one-time evaluations, not ongoing screening. It produces a file at the end — with no place to store a case, no history, no audit trail. That candor is a sign Anthropic understands exactly what it built.

Strip the KYC Screener to its essence and you have two genuinely useful primitives: **skills**, which package expertise into reusable instructions, and **data integrations**, which let an agent reach outside services. Both are necessary. Neither one, nor the two together, is a compliance program.

“


They shipped an agent that can use a compliance system. They didn't ship the compliance system itself. You still need a deterministic, rules-based system — because for every program we administer in every jurisdiction, we have to know exactly which documents we need, exactly which data we need, and exactly which watchlists to screen against. And it has to work the same way every time.

---

Kevin De Regt

VP of Product, Fund Administration, Juniper Square

- *This is the engine. The rest of this paper is about the car.*



#### THE REST OF THE CAR

# The parts required long before agents

Set AI aside. These components were required long before agents existed and will be required long after the novelty wears off. An engine that cannot connect to them moves nothing.

---

## STRETCH ONE

*When an investor arrives.*

The first stretch of the car is everything that happens when an investor arrives. You verify identity and run customer due diligence, escalating to enhanced due diligence when the risk calls for it. You assign a risk tier, because a US individual and a trust in a high-risk jurisdiction cannot be treated the same way. You screen every investor against sanctions, PEP, and adverse-media lists through controlled connections to providers such as LexisNexis Bridger or Refinitiv World-Check — rather than pasting names into a website. Hits come back with context, false positives get worked through a defined process, and an analyst confirms or clears each one with the reasoning recorded. You map beneficial ownership down to the threshold that applies — ten percent for a Cayman fund, twenty-five percent for a US entity — and check what was disclosed against registries. And you verify source of funds, and for higher-risk investors source of wealth, tracing the path from where the money came from to the commitment in front of you.

## STRETCH TWO

*After the investor is in —  
where the real work lives.*

The second stretch is everything that happens after the investor is in, which is where most of the real work lives and where a one-time tool stops being useful. Documents expire, so a program reviews investors on a recurring schedule and chases fresh materials before the old ones lapse. Circumstances change, and each change is its own trigger. A new payment instruction calls for a targeted review. A contribution over a threshold calls for a source-of-funds check. A move to a higher-risk jurisdiction reopens the file. A program that only knows how to onboard an investor is blind to all of this, and the obligations do not pause because the software was not watching.

## STRETCH THREE

*What makes any of it defensible.*

The third stretch is what makes any of it defensible. Material findings pass through maker-checker, where a second person reviews a determination before it stands and sign-off authority is tiered by risk. And every decision lands in an audit trail that survives examination, retained for seven years, capturing what was decided, on what information, by whom, and why. This is the part that matters most when a regulator arrives — because, as Joshua puts it, the regulator is not really investigating your answer.

“

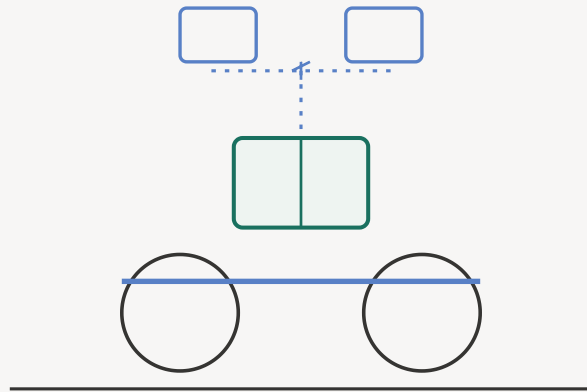
**Documenting and showing how you got there is more important than your actual outcome.**

— Joshua Demers

Head of Investor Services, Juniper Square

These are not enhancements you add once the basics are working. They are the basics. A GP who picks up an engine and tries to drive away will be a long way down the road before realizing the brakes, the steering, and the seats were never in the box.

- *The brakes, the steering, the seats — never in the box.*



## BUILDING THE CAR FOR THE AI-NATIVE ERA

# Eleven components. Cognition is one.

If you want to run a compliance program that genuinely uses agents — rather than demonstrating one — you need a specific set of components working together. A CFO or CCO can use this list to size up any provider that claims to be AI-native.

---

## THE ELEVEN, PART ONE

Everything in Part 2 has been true for years. What follows is newer, and it is the part a cookbook cannot hand you. Here is the list.

**You need cognition.** The processing power — reading and reasoning over messy documents and ambiguous cases. It is the one piece AI now does genuinely well, and the piece Anthropic shipped. It is also the easiest of these components to obtain, which is precisely why it is a poor place to plant a flag.

**You need accuracy and consistency.** Regulators expect the same case to produce the same result every time, and a probabilistic model on its own cannot promise that. Use the model only for judgment that genuinely requires fuzzy reasoning, and enforce the rules themselves in deterministic code — with continuous evaluation against real documents so drift is caught, not discovered.

**You need observability and auditability.** When a regulator walks in, they want proof you ran the program you described — which means seeing what your agents did, step by step, and reconstructing any case long after it closed. An agent whose reasoning evaporates when the session ends cannot meet this standard.

**You need tooling.** Agents have to operate the real instruments of the job: the screening services, the banking connections, the registries, the document stores, the general ledger. Each call has to run through a controlled, predictable interface rather than a loose request and a hopeful parse.

**You need data integrations.** The tools only matter if the agent can reach them — stable connections into the systems where investor records, fund structures, and positions already live, so the program operates on the real state of the business rather than a pasted-in snapshot.

**You need user workflows.** Collecting a driver's license from an LP, requesting a missing document, routing a question to a GP, moving a case through a queue. These are designed surfaces people use all day, and they are as much a part of the program as the model is.

## THE ELEVEN, PART TWO

**You need decision logic that is friendly to the LP.** The way a program asks for information shapes whether the investor stays patient or grows frustrated. A program designed around the investor's experience, not only the compliance officer's checklist, is a competitive advantage in its own right.

**You need human oversight.** Agents are confident even when wrong, so a program keeps people in the loop for the edge cases and for supervising the agents themselves. The right model is bounded autonomy: let the agent do the high-volume intake, triage, and preparation, and keep human accountability where the stakes are highest.

**You need agent memory and knowledge management.** An agent that forgets everything between sessions repeats work and loses context. The program needs a durable place for what the agents know and have seen, kept distinct from the system of record and managed deliberately.

**You need multi-user access.** In a real operation, many people work alongside many agents at once, across many clients. Permissions decide who can see which client's data, and the whole arrangement has to make sense as teams and agents grow.

**And you need a controls gateway.** An agent that can write to an investor profile is, by default, an unmonitored process pointed at production data. A controls gateway gives that agent a real identity, tied to a human's authority, and governs what it is allowed to touch — with circuit breakers that stop it when something looks wrong and the ability to roll back an action it should not have taken.

**1**

component is cognition — the one already widely available.

**10**

are where a compliance program is actually built, and where a decade of work either exists or does not.


**11**

components a buyer can walk, one by one, and ask a provider to show.

THE WHOLE VEHICLE

# An engine, and the car around it

Cognition is the engine — impressive, and the one part a buyer can simply purchase. Everything else is the program required to actually drive.



— Engine — cognition, the part you can buy    — The rest of the program — eight non-negotiables + the AI-native components

-- Drivetrain & chassis the engine bolts into

## THE RELATIONSHIP

# Why the LP experience is the whole game



## REASON ONE

*Don't go to jail. Obvious — and not the one firms underweight.*

## \$3.8B

global AML fines in 2025; a single TD Bank action ran ~\$3B on its own.

FENERGO, 2025 · BANKING DIVE, OCT 2024

There are two reasons a compliance program has to be excellent. The first is that you do not want to go to jail, or pay the kind of penalty that follows a real failure. That reason is obvious, and it is not the one firms underweight.

The second reason is the LP experience, and after staying out of trouble it is the most important thing a compliance program does. Onboarding is usually the very first workflow a new investor goes through with a GP. It is the first transaction, the first real interaction, the first impression — formed at the exact moment the investor is handing over a large sum of money and a great deal of sensitive personal information. If that experience is clumsy, the relationship starts on its back foot. When 87% of institutional investors have declined or reconsidered a commitment over AML/KYC concerns, the program is not a back-office function. It is part of the front-office relationship, and part of how a GP raises and keeps capital.

The trouble is that a good experience here is genuinely hard to deliver, and agents can make it worse before they make it better. There is a recursive quality to compliance: the program asks for a document, the document raises a question, the question prompts another request, and each loop is more friction for the investor. Designing that loop so it feels considered — asking for what is genuinely needed and no more — is the difference between a program that earns trust and one that erodes it.

## A NETWORK PROBLEM

This is where Juniper Square is unlike anyone else, because the experience problem is, at bottom, a network problem. An LP who has already been verified on the network does not start over with the next GP they invest with. Their compliance profile travels with them, reused across GPs on Juniper Square — with appropriate consent and a refresh of anything that has aged.

For the investor, that turns a dreaded process into one that mostly takes care of itself. For the GP, it turns compliance into a reason an LP is glad to be on the platform rather than a reason to grumble. No standalone agent, and no single-GP tool, can offer that — because none of them sit at the center of the network where the profiles already live.

---

**750K**

unique LP profiles on the platform

**1M+**

positions held across the network

**\$1T+**

active LP capital on Juniper Square

- *The experience problem is, at bottom, a network problem.*

## HOW TO EVALUATE A PROVIDER

# The buyer's view



## TWO WORDS

| [Show me.](#)

The eleven components in Part 3 are not only a build plan. They are a buyer's rubric, and a CFO or CCO evaluating any AI-native compliance provider can walk the list and ask, for each one, *show me*.

Show me where the rules live and prove they run the same way every time. Show me the evaluations against documents that look like mine. Show me the audit trail an examiner could read three years from now. Show me what stops an agent from taking an action it should not, and the path it follows when it is not sure. Show me what the LP actually experiences the second time they invest.

A provider who can answer clearly for each component is building a car. A provider who waves at the engine and skips the rest is selling you horsepower and leaving you to find your own wheels. We have turned that rubric into a one-page checklist — overleaf — so the comparison is easy to run against any vendor in a live evaluation.

- *One box checked at the top. Ten empty rows beneath it.*

## SCORED COMPARISON

# What an AI-native compliance program actually requires


Ask any provider to show you each one. An engine checks the first row — and only the first row.

		JUNIPER SQUARE THE WHOLE CAR	AN ENGINE ALONE COGNITION ONLY
01	<b>Cognition</b> <span>TABLE STAKES</span> Reads documents and reasons over messy, ambiguous cases.	✓	✓
02	<b>Accuracy and consistency</b> Same case, same result, every time; rules enforced in code, not guessed by a model.	✓	<input type="checkbox"/>
03	<b>Observability and auditability</b> Every step is visible and reconstructable for a regulator years later.	✓	<input type="checkbox"/>
04	<b>Tooling</b> Operates the real instruments: screening, banking, registries, the ledger.	✓	<input type="checkbox"/>
05	<b>Data integrations</b> Stable connections into the systems where investor and fund data already live.	✓	<input type="checkbox"/>
06	<b>User workflows</b> Designed surfaces for collecting documents, requesting items, and routing approvals.	✓	<input type="checkbox"/>
07	<b>LP-friendly decision logic</b> Asks for what is needed and no more; built around the investor's experience.	✓	<input type="checkbox"/>
08	<b>Human oversight</b> People in the loop for edge cases and for supervising the agents.	✓	<input type="checkbox"/>
09	<b>Agent memory and knowledge management</b> A durable, governed place for what the agents know.	✓	<input type="checkbox"/>
10	<b>Multi-user access</b> Many people and many agents working coherently across many clients.	✓	<input type="checkbox"/>
11	<b>Controls gateway</b> Governed agent identity, with circuit breakers and rollback.	✓	<input type="checkbox"/>

A standalone version of this checklist ships alongside the paper for live vendor evaluations. Cognition is the starting line, not the finish.

## THE CLOSING ARGUMENT

# Where Juniper Square stands



## THE GROUND IS SHIFTING

*FinCEN now treats advisers as financial institutions. Deadline: January 2028.*

The regulatory ground is shifting under every GP at once. In September 2024, FinCEN finalized a rule treating investment advisers as financial institutions under the Bank Secrecy Act, extending formal AML obligations to registered and exempt reporting advisers for the first time, with a compliance deadline now set for January 2028. As one law firm put it, these are obligations finally imposed on private fund managers.

At the same time, the volume is climbing fast. Preqin projects global alternatives AUM will reach roughly \$29 trillion by 2029, and as GPs raise more from individuals through the wealth channel, every new investor needs the same complete review as the institution that has been in the fund for a decade. The obligation is growing, the volume is growing, and the cost of getting it wrong is growing with them.

GPs have already decided how they want to handle this. Most of them want it off their plate, and they want one partner who can do all of it rather than a kit they have to assemble themselves.

## 91%

of GPs already outsource some or all of their AML/KYC work.

## 83%

name outsourcing to a service provider as their top priority for improving compliance —

## 59%

ahead of investing in new technology.

CSC Global, 2025

“

Tasks like KYC and AML are undifferentiated heavy lifting that GPs should hand to trusted partners rather than manage internally — so they can focus on their core competency of investing.

Alex Robinson  
CEO, Juniper Square

#### WHAT THE PARTNER BRINGS

What that partner needs to bring already exists at Juniper Square. The Trusted Human Services team — more than 425 fund administration professionals — runs real compliance programs today across onboarding, screening, refresh cycles, change-in-circumstance reviews, and audit discipline. Underneath them sits Agent Stack by Juniper Square: the knowledge harness, the private markets ontology and its native graph of investors, funds, and entities, and the orchestration for long-running work — with platform controls, agent identity, the controls gateway, and seven-year audit retention already in production.

The document-extraction AI that resolves offline subscription PDFs into structured fields already saves Joshua's team roughly 60% over manual entry. That is what AI looks like when it does the job it is suited for, inside a program that governs it.

## 425+

fund admin professionals running programs today

## ~60%

saved on document extraction over manual entry

## 40K+

investment entities on the network

THE ONLY  
CONFIGURATION THAT  
DRIVES

The compliance agent now in development runs on that foundation. It takes on document extraction, identity resolution, screening orchestration, source-of-funds preparation, and risk aggregation — working alongside the team under the same controls, the same audit trail, the same maker-checker, and the same continuous evaluation against real documents. The people stay in the seat of judgment, handling the high-risk cases and the relationships, while the agent absorbs the volume. That is the engine and the car together, which is the only configuration that actually drives.

Only Juniper Square can bring these pieces together at the velocity that AI now demands, on the secure, global, enterprise-grade platform that regulated GPs require, over an LP-GP operating system battle-tested for more than twelve years by more than 2,000 GPs. Anthropic built a remarkable engine, and the models underneath it get better every month. The car is the harder thing, and it is the thing a GP actually drives.

---

**Only Juniper Square can usher GPs into the AI era safely, swiftly, efficiently, and ahead of the pack.**

## SOURCES

CSC Global AML/KYC Report (November 2025; 2025). FinCEN final rule on investment advisers (Federal Register, September 4, 2024) and effective-date postponement (2025). Akin Gump client alert (September 2024). Preqin global alternatives AUM forecast (September 2024). Fenengo Global AML Fines Research Report 2025. Banking Dive reporting on the TD Bank enforcement action (October 2024). Internal sources: Juniper Square product review of Anthropic's KYC Screener; interview with Joshua Demers and Kevin De Regt.

---

**ABOUT JUNIPER SQUARE**

Juniper Square delivers the connected technology, data, and fund administration services that private markets GPs need to scale at any stage. The platform supports 40,000+ investment entities, 750,000 LP profiles, more than a million positions, and over \$1 trillion in active LP capital, and more than 2,000 GPs run on it worldwide. Juniper Square's 425+ fund administration professionals run compliance programs, fund administration, and investor services on behalf of GPs. Agent Stack by Juniper Square — which combines the knowledge harness, the private markets ontology, and agent orchestration — is the foundation on which Juniper Square's family of agents operates.