


Autonomous AI Trading, Constrained by Design

A technical and academic summary of the platform, investment process, portfolio-construction methodology, and deterministic risk controls.



Illustrative platform interface. Values shown are not presented as investment performance.

EDITION	AUTHOR	PLATFORM
4.0 · September 2026	Akan Udoeyop akan@vibeoscloud.com	VibeOS trade.vibeoscloud.com

This document describes software architecture and current operating methodology. It is informational and does not constitute investment advice, an offer, or a solicitation. Automated trading can result in loss of principal.

ABSTRACT

Platform summary and scope

Vibe Trader is an autonomous portfolio-management application for liquid U.S. equities and exchange-traded funds. It combines systematic market screening, a language-model decision layer, target-weight portfolio construction, deterministic risk enforcement, broker execution, and decision-level attribution. The system evaluates a curated opportunity set once per minute during its configured operating window. The model receives structured market, news, portfolio, and risk information and proposes a desired portfolio. Application code then determines whether each security is eligible, limits accepted weights, preserves required cash, reconciles working orders, and translates approved differences into broker instructions.

The platform is designed around a separation of functions. The model interprets evidence and states portfolio intent. It does not hold brokerage credentials, transmit orders directly, or alter risk parameters. The enforcement layer applies the mandate independently of the model response. This arrangement does not establish that the model will forecast returns accurately. It limits the authority attached to an uncertain forecast and records the basis, controls, execution status, and outcome of each decision.

SCOPE

This paper explains the current platform and investment process. It does not report an audited track record, a backtest, or an expected return. Screenshots are product documentation rather than performance evidence.

Keywords: autonomous trading; portfolio construction; target weights; trend following; risk budgeting; execution controls; decision attribution; model governance.

CONTENTS

Structure of the paper

SECTION	SUBJECT	PAGE
1	Abstract and scope	2
2	Platform definition and components	4–8
3	Operating thesis and control architecture	9–10
4	Opportunity set and market data	11–13
5	Signal methodology	14–15
6	Portfolio construction	16–18
7	Risk and exit methodology	18–20
8	Execution and operating model	21–22
9	Records, measurement, and advisor	22–28
10	Security, governance, and limitations	28–30

Terminology

Throughout this paper, model refers to the language model that proposes portfolio targets; engine refers to deterministic application logic; agent refers to the configured combination of model, mandate, account scope, and operating state; and broker refers to the external custody and execution venue.

1 · PLATFORM DEFINITION

What Vibe Trader is

An AI-powered trading platform that manages a brokerage portfolio autonomously, within limits that software enforces.

Vibe Trader is an agentic trading platform. Each user configures one or more AI agents, each backed by a large language model (currently Gemini 3.1 Pro) and assigned a mandate: a risk profile, an eligible universe, exclusions, a cash reserve, and a paper or live brokerage account. Once activated, an agent runs without user prompting. Every minute during its operating window it observes the market and its own portfolio, forms a view, decides what the portfolio should hold, and has the resulting orders placed with the broker. It then records what it saw, what it decided, and why.

The AI model is the agent’s judgment. It reads ranked market candidates, trend diagnostics, scored news, current holdings, cash, open orders, and the agent’s risk state, and it returns a complete target portfolio: which securities to hold, at what weights, and a written rationale for each change. It chooses among qualified opportunities, decides when a position’s thesis has broken, and decides when the book should be defensive and park cash in SGOV or hedge with GLD. The model does not screen the universe, calculate indicators, size orders, hold broker credentials, or override stops. Those functions belong to deterministic software.

What the platform ultimately does is convert that judgment into a controlled, auditable trading process. A systematic screen evaluates the full liquid universe; the model forms the portfolio view; a risk engine accepts, reduces, or rejects each requested weight under the mandate; an execution layer turns accepted differences into priced broker orders; and the decision record retains the full sequence, including holds and rejections. The user supervises rather than operates: they set the mandate, enable live trading, review decisions, and can pause any agent at any time.

AGENTIC CAPABILITY	WHAT THE AGENT DOES
Continuous operation	Evaluates the market and its book every minute without being prompted
Portfolio judgment	Proposes a full target-weight portfolio with a stated rationale
Event interpretation	Reads scored news and weighs fresh catalysts against price evidence
Position management	Adds, trims, or exits holdings as evidence changes; names the reason for each exit
Defensive posture	Chooses to park idle cash or hedge when no setup justifies risk
Self-documentation	Records context, proposal, controls, orders, and outcomes for every cycle

1 · PLATFORM DEFINITION

A portfolio-management system with bounded autonomous execution

Vibe Trader is organized around the account rather than around isolated trade recommendations.

The platform connects a user-controlled brokerage account to an automated decision process. Each account maintains separate credentials and a selected paper or live mode. An agent operates within that account scope and receives attributed holdings, available cash, open orders, market data, news, and its current risk state. The output of each evaluation is a proposed set of target weights rather than an unstructured series of buy and sell instructions.

The primary user surfaces are the portfolio view, order entry and history, agent administration, the decision record, market news, and the portfolio-aware advisor. The portfolio view reconciles broker-reported equity, cash, buying power, positions, and working orders. Agent pages separate attributed holdings and results from unrelated account activity. The decision record retains executed, skipped, and failed actions with the reason recorded at the time.

COMPONENT	FUNCTION	PRIMARY RECORD
Portfolio	Account value, positions, cash, orders, benchmark views	Broker snapshots
Trading	Manual market or limit order entry and order history	Broker orders and fills
Agents	Configuration, activation, holdings, and performance	Agent mandate and attributed fills
Decisions	Per-cycle proposals, controls, actions, and outcomes	Structured decision records
Market news	Normalized headlines, symbol mapping, sentiment, and impact	Scored news items
Advisor	Account-aware analysis and user-confirmed order proposals	Conversation and cited account context

1 · PLATFORM DEFINITION

Account, agent, and brokerage boundaries

A Vibe Trader account is an operational boundary. Positions, agents, decisions, market mode, and credentials are scoped to the selected account. Shared accounts assign owner, trader, or viewer permissions. Platform-level roles—administrator, trader, and read-only—remain separate from account membership and cannot be expanded by a workspace invitation.

Paper and live brokerage environments are treated as distinct books. Credentials, positions, orders, and histories are selected by mode. Live automation requires saved live credentials and an explicit

live-trading enablement setting. A paper agent cannot access the live book, and a live agent does not infer positions from the paper account.

BOUNDARY	INCLUDED	EXCLUDED
Account	Members, agents, decisions, credentials, settings	Other accounts
Agent ownership	Lots created by that agent’s filled orders	Manual lots and other agents’ lots
Broker mode	Paper or live positions, orders, and history	The alternate mode
Role	Operations authorized for the member	Platform privileges not separately granted

Attribution Rule

Autonomous liquidation authority follows fill attribution. An agent may manage positions assigned to its own book; it does not treat every security visible in the brokerage account as its property.




Figure 1. Portfolio and brokerage state. The portfolio workspace presents broker equity, buying power, portfolio and benchmark paths, positions, orders, and watch data in one account-scoped view. The application is shown in its onyx dark theme; displayed values are illustrative product documentation.




Figure 2. Agent supervision and attributed performance. The agent workspace combines operating status, risk profile, attributed results, benchmark comparison, holdings, and recent decisions. Controls remain separate from the model's portfolio proposal.

Edit agent X

Pick a model, a risk profile and the hard caps the agent can never breach.

Name

Model: Gemini 3.1 Pro — Deeper reasoning, slower

Risk tolerance: Balanced

Trading mode: Paper

End-of-day policy: Flatten losers

Agent picks its own universe ☒

Finds candidates from market news, live movers, and the session's most active names.

Discovery source: News + movers (recommended)

Movers are only discovery signals, the agent still confirms liquidity, trend, and catalyst context before buying.

Universe: AAPL, MSFT, NVDA, AMD, TSLA, AMZN, GOOGL, META, SPX, QQQ, TQQQ

Not used while the agent picks its own symbols.

Never trade: TSLA

Agent sizes its own trades ☒

Derives max per position and total spend cap from live account equity and its risk tolerance.

Cash reserve from brokerage account ☒

Uses the real account balance to set the reserve instead of a fixed dollar amount.

Use entire portfolio balance ☒

The whole account is deployable, no cash is held back, and the agent decides each trade size from the remaining available balance.

Max per position (\$): 5000

Total spend cap (\$): 25000

Managed by the agent

Cash reserve (\$): 40000

Max trades per day: 1000

Managed by the agent

Drawdown stop (%): 12

Fractional shares ☒

Allows dollar-based sizing on market orders.

Extended hours ☒

Lets the agent place and cancel limit orders outside the session.

Allowed order types

Figure 3. Agent mandate configuration. The configuration surface defines risk, cash, position, loss, and operating constraints. These settings bound the model's authority and are applied by deterministic software.

2 · OPERATING THESIS

Forecasting judgment and portfolio authority are separate functions

The operating thesis is limited and testable: a language model may be useful for interpreting heterogeneous evidence and expressing a coherent portfolio view, but it should not be the final authority over eligibility, sizing, order mechanics, or loss controls. Vibe Trader therefore places the model between data preparation and policy enforcement.

The model receives normalized evidence rather than an unrestricted data stream. It is asked to identify the securities it wishes to hold, assign target weights, and state a rationale. The engine then evaluates that proposal under the current mandate. An accepted target may equal the request, be reduced, be replaced by a mandatory exit, or be rejected. The broker subsequently determines whether the accepted instruction can be filled.

01 Observation Quotes, bars, trends, news, positions, cash, and open orders.	02 Judgment The model proposes a complete target portfolio and rationale.	03 Authority The engine applies eligibility, size, cash, risk, and execution rules.
04 Instruction Accepted target differences become broker orders.	05 Outcome Acknowledgements, fills, errors, and P/L are recorded.	06 Review Results are available for operational and investment-process analysis.

This separation controls authority; it does not eliminate model error. A model can still rank valid securities poorly, misinterpret news, or select a portfolio that underperforms. Deterministic controls bound specified consequences but do not convert a forecast into certainty.

2 · OPERATING THESIS

Mandate hierarchy and decision precedence

LEVEL	DECISION OWNER	EXAMPLES
1	Account owner	Broker mode, live enablement, exclusions, reserve, activity limits
2	Model	Desired securities, target weights, current rationale, defensive posture
3	Policy and risk engine	Eligibility, concentration, turnover, stops, daily states, ownership
4	Broker	Tradability, available quantity, buying power, acceptance, fills

The hierarchy resolves apparent conflicts. A model-requested defensive allocation to SGOV or GLD may be retained when it states a specific market-risk rationale, because portfolio posture is part of the model's assigned judgment. The same model cannot exceed the account's symbol cap, use an excluded security, waive a stop breach, or sell shares that are not available. Those matters belong to the mandate, enforcement, or broker layers.

The engine can also produce a position that was not explicitly selected by the model when a configured exposure floor is active and ranked candidates satisfy the entry gate. Such positions are marked as engine-owned and receive additional preservation rules. This behavior is disclosed in the record because it changes the source of portfolio intent.

**CONTROL
INTERPRETATION**

"Constrained by design" means that authority is partitioned and independently evaluated. It does not mean that losses are prevented or that every control will perform as intended under discontinuous markets.

3 • OPPORTUNITY SET

A curated liquid universe evaluated in full each minute

The current opportunity set contains 187 U.S.-listed equities and funds selected for liquidity and relevance to the strategy. Each active agent evaluates the full list on every scheduled cycle, subject to its configured universe and exclusions. Membership is a data-collection permission, not a trade recommendation.

STAGE	TEST	RESULT
Curated pool	187 listed equities and funds	Available for market-data refresh
Agent mandate	Configured universe less exclusions	Permitted for that agent
Price and liquidity	Price $\geq$ \$5; average daily dollar volume $\geq$ \$50 million	Removes structurally difficult instruments
Transaction cost	Quoted spread $\leq$ 0.30%	Limits expected entry friction
Movement	ATR at least 0.60%, or clear intraday range expansion on above-pace volume	Requires sufficient observable range
Signal and portfolio fit	Entry path, extension, regime, cash, sector, ownership	Candidate may enter target construction

ROLE OF THE AI MODEL

The model does not choose or screen the universe. Screening is deterministic and applied to all 187 names every cycle. The model sees only held securities and the highest-ranked candidates that passed the screen, and it may add a name to its portfolio only from that set.

Leveraged and inverse products are excluded from the standard liquid universe. SGOV and GLD are included for defined capital-sleeve functions rather than treated as ordinary trend positions. The effective opportunity set can be smaller when data are missing, spreads widen, a symbol is excluded, or a same-session stop makes a security temporarily ineligible.

3 • OPPORTUNITY SET

Market, portfolio, and event data

A decision cycle uses several observation horizons because no single statistic defines the investment case. Intraday measures describe the current leg; daily measures describe recent and structural trend; relative measures compare the security with SPY; event data supply a possible reason for repricing; and portfolio data determine whether the idea can be owned at an acceptable size.

DOMAIN	VARIABLES	PURPOSE
Intraday	Session change, VWAP, 15-minute slope, acceleration, range location, 60-minute high, paced volume	Current direction and participation
Medium horizon	5- and 20-day return, SMA20 distance, RSI(14), ATR(14)	Swing trend, extension, and volatility
Long horizon	SMA50, SMA200, 52-week range	Structural trend context
Relative	5- and 20-day return less SPY return	Leadership against the benchmark
Event	Headline age, symbols, sentiment, impact, assessment	Potential catalyst and risk context
Portfolio	Cash, positions, cost, peak, stops, open orders, day P/L	Capacity and current risk

MISSING DATA

An unavailable indicator is represented as unavailable. It is not interpreted as negative evidence. This distinction prevents a data failure from becoming a bearish signal.

4 · DECISION PROCESS

The per-minute evaluation cycle

Active agents are scheduled from 30 minutes before the regular-session open through 30 minutes after the close. The cycle is designed to refresh the state of the opportunity set and portfolio rather than to compel a trade. A hold or no-order result is valid when the accepted target book already matches the attributed book or when no candidate satisfies the gate.

STEP	OPERATION	OUTPUT
1	Refresh broker account, positions, open orders, and market clock	Current executable state
2	Fetch and calculate data for the full liquid universe	Security-level observations
3	Apply price, liquidity, spread, trend, extension, and regime gates	Eligible candidates
4	Rank candidates and combine them with attributed holdings and news	Compact decision context
5	Request a complete target-weight portfolio and rationale	Model proposal
6	Apply risk, ownership, cash, concentration, turnover, and exit rules	Accepted target book
7	Calculate order deltas and reconcile working orders	Executable instructions
8	Persist the context, proposal, controls, order status, and outcome	Decision record

ROLE OF THE AI MODEL

The model is responsible for step 5 only: it reads the compact decision context and returns a complete target portfolio with a rationale for each change. Steps 1–4 and 6–8 are performed by software, so a model failure or timeout leaves the existing book and its stops in place rather than producing unmanaged orders.

The full-universe refresh is separate from the model context. The engine may calculate observations for all 187 securities while sending the model a compact set consisting of current holdings and the highest-ranked valid candidates. This limits irrelevant prompt volume without reducing the breadth of systematic screening.

5 · SIGNAL METHODOLOGY

Five entry paths share one eligibility framework

Candidates can enter consideration through established trend, emerging trend, catalyst confirmation, intraday leadership, or intraday reclaim. The paths describe different market states; they do not bypass common liquidity, spread, extension, portfolio, or regime controls.

PATH	EVIDENCE	INTERPRETATION
Established trend	Price above SMA20 and SMA50; positive 20-day return and relative strength	Persistent leadership
Emerging trend	Positive session move and 30-minute momentum; upper-range location; adequate volume; positive 5-day relative strength	Early session participation
Catalyst	Fresh high-impact event plus positive price confirmation and acceptable range position	Event-associated repricing
Intraday leader	Above VWAP for at least ten minutes; positive 15-minute slope; upper third of the session range; session return at least 0.5% and 0.5 points above SPY; paced volume at or above normal	Sustained session leadership
VWAP reclaim	Above VWAP for at least three minutes; positive 15-minute ATR-normalized slope; acceleration; intact longer trend	Recovery after an earlier decline

The reclaim path addresses a limitation of open-to-current returns. A security can remain negative on the session while appreciating materially from a later low. Requiring a sustained move above VWAP, positive recent travel, acceleration, and supporting trend context distinguishes a current recovery from a single favorable print. It remains a candidate test rather than an automatic buy rule. The leader path serves a related purpose: a security that climbs steadily through the session can fail the momentum thresholds of the emerging path on any single minute while its sustained position above VWAP and outperformance of SPY are unambiguous.

ROLE OF THE AI MODEL

Qualification is decided by software; the model cannot promote a security that failed the gate. The model decides which qualified candidates to own and in what proportion, using the setup label, ranking diagnostics, and any related news to judge which opportunities are strongest and whether they fit the current book.

5 · SIGNAL METHODOLOGY

Candidate ranking and extension controls

Eligibility and priority are separate. The gate determines whether a security may be considered. A composite score then orders valid candidates so that current leadership can be compared with medium-horizon and benchmark-relative evidence.

COMPOSITE SCORE

$$S = 0.50 \cdot RS20 + 0.75 \cdot RS5 + 0.30 \cdot M20 + 1.50 \cdot Z50 + 1.25 \cdot V + 0.50 \cdot M30 + 3.00 \cdot L15 + 1.50 \cdot D_{VWAP} + A + H$$

RS denotes relative return versus SPY; M denotes medium or intraday momentum; Z50 is ATR-normalized distance above SMA50; V is paced volume; L15 is recent ATR-normalized slope; D_{VWAP} is distance from session VWAP; A is acceleration; and H captures proximity to a 60-minute high. Terms are bounded where appropriate before combination.

Recent slope receives the largest explicit coefficient because the strategy is intended to distinguish a security rising now from one whose earlier gain is fading. Structural and relative-strength terms reduce reliance on the most recent interval. RSI at or above 80 or a price more than approximately three ATR above SMA20 is treated as excessive extension and blocks a new entry.

OBSERVATION	EFFECT
Positive recent slope and acceleration	Raises priority
Persistent position above VWAP	Raises priority
Near a 60-minute high	Adds local-breakout confirmation
Lower part of the session range without a catalyst	Rejects fading tape
Extreme RSI or ATR extension	Blocks entry

6 · PORTFOLIO CONSTRUCTION

Target weights define the desired portfolio state

Vibe Trader does not treat each signal as an independent bet. The model returns a target portfolio. For security i , the target weight is the desired market value divided by the deployable portfolio base. The base includes attributed positions marked to market and cash available after reserves and working orders.

TARGET WEIGHT

$$w_i^* = \text{desired market value}_i / \text{deployable portfolio base}$$

ORDER DELTA

$$\Delta_i = (w_i^* - w_i) \times \text{deployable portfolio base}$$

Reductions are evaluated before increases so that sales can fund stronger candidates in the same cycle. Differences inside the no-trade band are ignored. Fresh working buys count toward the target already committed; fresh working sells count as pending reductions. This prevents repeated orders from being created while the broker is still processing the prior instruction.

ROLE OF THE AI MODEL

The model is responsible for the desired portfolio: the set of securities and the target weight of each. The engine is responsible for everything that converts that intent into an accepted book: caps, cash, sizing, and order deltas.

A target-weight representation makes opportunity cost observable. Raising one allocation affects cash, sector exposure, and room for other names. It also permits the engine to record requested weight, accepted weight, the reason for any clamp, and the resulting order rather than recording only the final transaction.

6 · PORTFOLIO CONSTRUCTION

Volatility-normalized sizing and allocation limits

Position size is bounded by both portfolio allocation and planned loss at the initial stop. The current risk budget is 0.50% of account equity per new position. Initial stop distance is the greater of 0.50% or ATR percentage multiplied by the risk-profile factor: 1.25 for conservative, 1.50 for balanced, and 1.75 for aggressive.

RISK-SIZED NOTIONAL

$$\text{Notional} = \min(\text{allocation cap}, (\text{Equity} \times 0.50\%) / \text{stop distance } \%)$$

ILLUSTRATIVE INPUT	CALCULATION
Account equity	\$100,000
ATR percentage	2.0%
Balanced stop distance	$1.50 \times 2.0\% = 3.0\%$
Dollar risk budget	$0.50\% \times \$100,000 = \$500$
Risk-sized notional	$\$500 / 3.0\% = \$16,667$
Starter allocation cap	$5.0\% \times \$100,000 = \$5,000$
Initial accepted notional	\$5,000

The example is mechanical rather than predictive. The accepted amount can be lower because of cash reserves, available buying power, symbol and sector limits, existing exposure, daily risk state, working orders, or broker restrictions.

6 • PORTFOLIO CONSTRUCTION

Portfolio shape and accepted-weight example

CONTROL	CURRENT DEFAULT	PURPOSE
Starter position	5%	Limits initial cost of a false signal
Standard slot	10%	Provides economically meaningful exposure
Single-name ceiling	5%–15%, configuration-dependent	Limits idiosyncratic concentration
Concurrent positions	Up to 8	Balances breadth and monitoring quality
Sector ceiling	55% of capital	Limits clustered sector error
Risk-on exposure range	60%–80% when valid candidates exist	Maintains participation without requiring full investment
Cash reserve	Owner-configured	Retains unavailable capital

The single-name limit is calculated from the configured maximum position value and deployable capital, with an absolute ceiling of 15%. Exposure is not assigned to securities that fail the entry gate. If the accepted model portfolio is below the risk-on floor and valid candidates exist, the engine may add ranked starter positions. A specifically justified model request for defensive parking takes precedence over that backfill.

SECURITY	CURRENT	PROPOSED	ACCEPTED	RESULT
AAPL	10%	10%	10%	No order

SECURITY	CURRENT	PROPOSED	ACCEPTED	RESULT
PLTR	12%	15%	12%	Capped by account setting
HOOD	0%	8%	5%	Starter buy
NVDA	8%	0%	4%	Deep exit reduced without thesis break
SGOV	15%	0%	0%	Released to fund accepted risk
Cash	55%	—	69%	Residual

7 · RISK METHODOLOGY

Position risk changes as a trade progresses

The exit framework is asymmetric. Initial downside is defined before entry. As unrealized profit increases, the position moves through a set of risk states. These transitions are computed from the entry price, initial risk unit, peak price, current price, and ATR; the model receives the resulting state but cannot waive a deterministic stop breach.

STATE	CONTROL	INTERPRETATION
Initial	Entry less $k \times \text{ATR}$	Full planned downside remains
+0.5R	Stop advances to entry	A meaningful winner is protected from becoming a loss
+1R	Approximately $2 \times \text{ATR}$ trailing stop	Fixed target is removed; trend may continue
+2R	Eligible for partial realization; approximately $3 \times \text{ATR}$ trail	Some profit may be realized while remaining exposure receives a wider leash
Stop breach	Mandatory exit	Model preference does not supersede the breach

R is the initial dollar risk per share: entry price less initial stop price. A wider trail at higher R-multiples is intended to reduce premature exits from normal volatility. It also permits a larger giveback from the peak. The trade-off is explicit: tighter protection realizes gains sooner but increases the probability of exiting a continuing trend.

7 · RISK METHODOLOGY

Thesis changes, stale capital, and re-entry

Elapsed time is treated as diagnostic information rather than an independent reason to sell. A profitable position above its stop is governed by the trailing framework. A meaningful loser remains subject to its predefined stop. A position that has been held for at least 90 minutes, has gained less than 0.3R, is not rising, and is not already a material loser may be recycled into a stronger qualified candidate.

CONDITION	TREATMENT
Working winner	Hold unless the trail, a named thesis change, or a mandatory control requires action
Material loser at or below $-0.35R$	Manage through the defined stop rather than a time limit
Rising from the lower third of the session range	Do not classify as stale while the recovery is developing
Stagnant for at least 90 minutes and below $+0.3R$	Eligible for capital recycling
Model requests a full exit without a named thesis change	Reduce the depth of the requested cut
Stopped out earlier in the session	Ineligible for re-entry for the remainder of the session

ROLE OF THE AI MODEL

The model receives each position's risk state and may request a reduction or exit when it names a specific change in the thesis. It cannot waive a stop breach, loosen a trail, or lift a daily risk state. A model-requested exit without a named reason is moderated rather than executed in full.

There is no general post-trade cooldown. A separate quality control applies after repeated losing exits. These rules are intended to reduce rapid churn and low-quality re-entry without imposing a blanket minimum holding period.

7 · RISK METHODOLOGY

Daily risk states and capital sleeves

Position stops address idiosyncratic errors. Daily states respond to correlated or repeated loss across the account. Under the standard configuration, a decline of 1.5% places the agent in an exits-only soft-stop state. At a daily gain of 2%, new entries are reduced to half size. The configured hard drawdown limit is bounded at no worse than 4%, after which weight increases are prohibited.

STATE	THRESHOLD	NEW RISK
Normal	Above -1.5% and below $+2.0\%$	Permitted within the mandate
Soft stop	At or below -1.5%	Exits and reductions only
Lock-in	At or above $+2.0\%$	New entry size halved
Hard halt	Configured threshold, capped at -4.0%	No weight increases

Cash management is distinct from equity selection. Excess cash may be allocated to SGOV when no equity setup qualifies. SGOV can be released when an accepted opportunity requires capital. GLD is an optional risk-off hedge capped at 8%; it is not used as a general cash substitute. SGOV and GLD are excluded from ordinary equity ATR stops, dead-money recycling, position-count limits, and turnover controls, while their explicit sleeve caps and available-cash constraints remain applicable.

RESIDUAL RISK

Daily thresholds do not guarantee a maximum realized loss. Gaps, partial fills, rejected orders, and rapid price changes can produce outcomes beyond the planned limit.

8 · EXECUTION

Order construction is part of the investment process

A valid signal can produce an unfavorable result if transaction costs or order handling are ignored. The execution layer therefore considers session status, quoted spread, available quantity, working orders, and broker restrictions. During the regular session, the platform uses marketable limits where applicable: buys are priced near the ask and sells near the bid, avoiding unrestricted price exposure while seeking prompt execution.

SESSION OR STATE	ORDER TREATMENT	RATIONALE
Regular session	Marketable limit where applicable	Bound price exposure while seeking execution
30 minutes before open	Extended-hours limit	Permit resting instruction in thinner liquidity
30 minutes after close	Extended-hours limit	Permit bounded after-hours execution
Outside window	No scheduled agent evaluation	Avoid unattended continuous operation
Fresh working buy	Count toward committed target weight	Prevent duplicate buying
Fresh working sell	Count as pending reduction	Prevent duplicate exits
Stale sell after five minutes	Cancel, refresh quantity, then reconsider	Release reserved shares before replacement

ROLE OF THE AI MODEL	None in order construction. The model never specifies prices, quantities, or order types and has no access to broker credentials. Execution is entirely deterministic.
----------------------	--

Broker acceptance is not a fill. The actual book can differ from the accepted target because an order remains open, fills partially, is rejected, or encounters insufficient available quantity. The decision record preserves these distinctions.

8 · EXECUTION

Ownership attribution and manual control

Manual and autonomous activity can coexist in one brokerage account, but they are not interchangeable for attribution. Agent lots are constructed from that agent’s filled orders. Manual holdings and positions attributed to another agent are excluded from its autonomous book. This prevents a target-weight reduction from automatically liquidating unrelated assets.

Users retain operational control. An agent can be paused or resumed, a manual evaluation can be initiated, and positions can be managed directly. Pausing stops new evaluations; it does not by itself close existing positions. Broker-side orders and positions remain visible through the portfolio and order views.

EVENT	ATTRIBUTION EFFECT
Agent buy fill	Creates or adds to that agent's FIFO lots
Agent sell fill	Reduces attributed lots and realizes P/L against their cost
Manual trade	Changes broker account but is not assigned to an agent automatically
Other-agent trade	Belongs to the originating agent's book
Broker fee or settlement adjustment	Included in account-level results; reconciled where attributable

PRACTICAL CONSEQUENCE

Account-level portfolio movement and agent P/L can answer different questions. Reconciliation requires the same timestamp, the same scope, and explicit treatment of cash, fees, manual activity, and open positions.

9 • RECORDS AND MEASUREMENT

A decision record includes inaction

Order history alone is insufficient to reconstruct an autonomous process. Vibe Trader stores the market and portfolio context available to the agent, the proposed target book and rationale, the controls applied, the resulting instruction, and the broker outcome. A hold, blocked action, or error is retained because it describes the behavior of the process under observed conditions.

RECORD LAYER	EXAMPLES
Context	Market regime, ranked candidates, quotes, trend diagnostics, news, holdings, cash, open orders
Proposal	Requested symbols, target weights, rationale, cited events, posture
Enforcement	Accepted weights, rejected securities, clamp reasons, turnover and daily state
Execution	Side, quantity or notional, order type, limit price, broker status, errors
Outcome	Fill time and price, realized sell P/L, current valuation
Inaction	No valid candidate, no executable delta, risk state, or precise eligibility failure

ROLE OF THE AI MODEL

The model authors the rationale stored with each decision: its stated interpretation of the evidence and the reason for each requested change. All other record layers are written by software.

Natural-language reasoning and structured fields serve different purposes. The rationale records the model's stated interpretation. Structured market, risk, and execution fields allow the statement to be compared with contemporaneous evidence and controls. Neither should be treated as proof that the investment thesis was correct.




Figure 4. Decision record and execution status. Each cycle retains the model rationale alongside accepted actions, skipped actions, execution state, and financial attribution. The record includes inaction because a hold or rejection is part of the investment process.

9 • RECORDS AND MEASUREMENT

Performance definitions and benchmark construction

Portfolio return and agent P/L are not interchangeable. Portfolio measures include all account assets and activity. Agent measures include attributed fills and marked open lots. Both require a defined start time and synchronized observations before comparison with a benchmark.

MEASURE	DEFINITION	SCOPE
Portfolio value	Broker-reported equity	Cash, positions, fees, and all account activity
Portfolio change	Change from the first plotted account value	Selected account and interval
Agent realized P/L	Sell proceeds less FIFO cost of attributed lots	Closed agent activity
Agent unrealized P/L	Current broker value less cost of open attributed lots	Open agent activity
Agent total P/L	Realized plus unrealized; broker net result may be used when the agent owns all activity	Agent contribution

MEASURE	DEFINITION	SCOPE
SPY value path	Opening portfolio value multiplied by SPY price relative	Counterfactual all-SPY path
SPY P/L path	Benchmark change rebased to zero at the first plotted bucket	Same selected interval

Broker equity, quotes, fills, and application snapshots update on different clocks. Temporary differences can therefore occur even when definitions are correct. Formal evaluation should use broker statements, synchronized observations, and explicit treatment of deposits, withdrawals, fees, and strategy changes.

**EVIDENCE
STANDARD**

The current paper presents methodology only. Any future performance presentation should identify the account, paper or live status, period, benchmark, cash flows, gross or net treatment, fees, slippage, material changes, and whether results were independently verified.

9 • RECORDS AND MEASUREMENT

Market intelligence and the portfolio-aware advisor

The market-intelligence layer normalizes news, maps affected symbols, and records sentiment, estimated impact, trend classification, and a concise assessment. The same processed information is used by the market-news interface, included selectively in agent context, and available to the advisor. News is one input to the process; it is not required for established, emerging, or reclaim trend paths.

The advisor is a conversational interface over the selected account’s portfolio, orders, attributed agent results, decision history, news, and SPY comparisons. It can answer questions about sources of profit and loss or recent actions using current account data. When it prepares an order, the order remains subject to user confirmation and the relevant execution controls.

01 Ingest Retrieve headlines and associated symbols.	02 Normalize Standardize source, time, mapping, and text.	03 Assess Record sentiment, impact, event summary, and watch condition.
04 Distribute Serve news, agent context, and advisor context.	05 Decide Combine events with market and portfolio evidence.	06 Record Retain cited events with the resulting decision.
ROLE OF THE AI MODEL AI models perform three functions here: assessing each headline for sentiment, impact, and a short summary; weighing assessed news inside the trading agent’s decision; and powering the advisor’s answers. The advisor cannot place orders on its own; any order it prepares requires user confirmation.		

The advisor supplies information and analysis, not individualized investment advice. Its account access improves relevance but does not establish the correctness of its conclusions.




Figure 5. Market-intelligence workspace. Normalized headlines, symbol mappings, sentiment, impact, and watch conditions provide a common event record for the news interface, trading agent, and advisor.




Figure 6. Portfolio-aware advisor. The advisor uses selected-account positions, orders, agent decisions, performance, and market intelligence to answer portfolio questions. Any proposed transaction remains subject to explicit user confirmation.

10 · GOVERNANCE

Security and permission model

Brokerage credentials are accepted through account settings, encrypted, and used only by server-side processes. Secret values are not returned to the browser after storage. User sessions are issued through the external sign-on system, are revocable, and can require multi-factor verification. Trusted-device state is signed and time-limited.

CONTROL AREA	IMPLEMENTATION INTENT
Credential handling	Keep broker secrets outside browser-readable application state
Account isolation	Scope queries and actions to authenticated account membership
Role separation	Distinguish platform role from account role
Live-trading gate	Require live credentials and explicit activation
Paper/live separation	Use separate broker environments, positions, and order histories
Auditability	Retain decisions, errors, orders, and role-relevant actions
Read-only access	Prevent settings changes, account linking, agent creation, and autonomous operation

Security controls reduce specified risks but do not eliminate compromise, software defects, third-party outages, or credential misuse. Institutional deployment should define key rotation, incident response, access review, evidence retention, and operational ownership beyond the application defaults.

10 · GOVERNANCE

Model and strategy change control

A model, prompt, signal, sizing rule, or risk constant can alter portfolio behavior and should therefore be treated as an investment-process change. The minimum governance record should identify the owner, effective date, intended effect, evidence reviewed, tests performed, release decision, and rollback point.

FUNCTION	MINIMUM CONTROL
Mandate owner	Approves risk posture, exclusions, reserve, broker mode, and live enablement
Strategy owner	Documents proposed model, signal, sizing, or prompt change
Independent review	Challenges assumptions, failure modes, data sufficiency, and evidence labels
Pre-release validation	Tests deterministic rules, boundary cases, and reconciliation behavior
Release gate	Requires a defined version, approver, effective date, and rollback path
Kill switch	Pauses agents or disables live trading without deleting history
Post-incident review	Reconstructs context, proposal, enforcement, broker response, and portfolio outcome

ROLE OF THE AI MODEL

The model itself is a governed component. Changing the model version or its instructions can change portfolio behavior as much as changing a risk constant, and should pass the same review and release gate.

The stored decision history can support research, but it does not make the system self-validating. Changes should not be approved solely because they would have improved a small number of recent observations. Evaluation should distinguish hypothesis formation from out-of-sample testing and account for regime dependence, transaction costs, and multiple comparisons.

10 · GOVERNANCE

Limitations, residual risks, and conclusion

Vibe Trader is an implementation of a bounded autonomous process. Its controls address defined operational and portfolio risks, but the system remains exposed to market uncertainty and to errors in models, data, software, execution, and human configuration.

RISK	LIMITATION
Market discontinuity	A price can gap through a stop; realized loss may exceed planned 1R.
Liquidity deterioration	A previously acceptable spread can widen and prevent execution.
Data latency or absence	Ranking and risk state can be stale; the engine may withhold new risk but cannot recreate missing information.
Model error	The model can misread evidence or rank valid securities poorly.
Correlation	A sector cap is a coarse control; securities can share hidden factors.
Regime lag	Benchmark-based state can react after the underlying market has changed.
Broker or network failure	Orders can be rejected, delayed, duplicated externally, or partially filled.
Short evidence history	Observed results may not support an inference of durable excess return.
Configuration risk	Owner-selected limits can be unsuitable for the account or objective.

The platform's central design decision is the separation of market interpretation from portfolio authority. A broad liquid universe is evaluated systematically; the model expresses a complete target portfolio; deterministic software applies the mandate and risk controls; the broker determines executable reality; and the record retains the full sequence. This structure makes the process inspectable and governable. It does not guarantee profitability, benchmark outperformance, capital preservation, or uninterrupted operation.

Disclosure. This paper describes the current software methodology as of September 2026. It is not investment advice, an offer, or a solicitation. Automated trading involves material risk, including loss of principal. Stops and limits may not execute at intended prices. Paper trading and hypothetical examples may differ materially from live results. Past performance, if later presented, is not indicative of future results.