TradeSense — User Handbook

> **Your autonomous, multi-agent AI trading platform.**
> Built for serious traders who want institutional-grade analysis without the institutional price tag.

Table of Contents

🚀 Quick Start — Initial Setup

To enable trading on your own account, you **must** configure your personal Alpaca keys in the dashboard.

1. **Open Settings:** Click the **Gear Icon (⚙️)** in the top right corner of the dashboard.
2. **How to get your Alpaca Keys (Step-by-Step):**
   - **Step 1:** Go to the Alpaca Website.
   - **Step 2:** Sign up specifically for the **"Trading API"** account.
   - **Step 3:** Verify your email address, then log in to your new account.
   - **Step 4: Enable MFA (Required for API access).**

     > [!IMPORTANT]
     > You MUST enable MFA before Alpaca will permit you to generate API keys.

- **Step 5:** Click the **Alpaca Logo** (top left) and change the toggle to **"Paper Trading"** (to use virtual/fake money).
- **Step 6:** In the side bar, click the **Home** button.
- **Step 7:** Scroll down to the bottom of the **right-hand side** page. You will find a section for **"API Keys"** and **"API Details"**.
- **Step 8:** Click "Generate New Key" (or "View"). Copy the **Key ID** and **Secret Key**.

3. **Enter in TradeSense:** Go back to the dashboard, click the ⚙️ Gear → Paste the keys → Click **Update Keys**.

> [!NOTE]
> The AI Brain (Cerebras, Groq, Gemini) is powered by **TradeSense Infrastructure** — you do **not** need to provide your own LLM keys. You only provide your broker keys so the AI knows where to send your approved trades.

> [!IMPORTANT]
> Without your own Alpaca keys, the "Approve" button will fail and live prices will not update correctly. The default keys in the source code are for demonstration only and are periodically rotated/disabled.

---

1. What is TradeSense?

TradeSense is a **10-agent LangGraph pipeline** that analyses financial assets using a combination of:

- **Mathematical indicators** (RSI, MACD, Bollinger Bands, ATR, EMA)
- **SEC Government EDGAR data** (real 10-K, 10-Q, 8-K filings you can actually read)
- **4-stream sentiment analysis** (Finnhub + Alpaca/Benzinga + StockTwits + Google RSS)
- **Self-hosted adversarial AI debate** (Bull vs Bear researchers argue before any trade happens)
- **Human-in-the-loop approval** (you are the final decision maker — no trade fires without your click)

All trades execute via **Alpaca Paper Trading** (free, $100,000 virtual balance). Zero real money at risk.

---

2. The Dashboard — Navigation

The top navigation bar has **7 main sections**:

Tab	What it does
Terminal	Run a manual AI analysis on any ticker. This is your main trading interface.
Debate	Read the full Bull vs Bear debate transcript from your last analysis.
Positions	View all currently open positions and their live P&L.

Tab	What it does
History	All closed trades — entry/exit prices, P&L, rationale.
Autopilot	Configure and control the autonomous trading engine.
Brain	Knowledge Base — SEC filings, learned lessons, macro data.
Backtest	Walk-forward historical simulation — prove your strategy with real data before risking capital.

The **top bar** also shows:

- **Live ticker tape** — Real-time prices for your tracked assets
- **AI Scanner button** — Opens the full 60-asset AI Scanner panel
- **Alpaca Paper balance** — Your current virtual portfolio value
- **Settings gear** — API key management

---

3. Terminal Tab — Manual Analysis

This is where you run on-demand analysis.

How to use it:

1. **Type a ticker symbol** in the search box (e.g., `AAPL`, `BTC`, `RELIANCE`, `EURUSD`)
2. Click **Run Analysis**
3. Watch the **9 AI Agents** light up one by one in real-time as they process
4. When all Phase 1 & 2 agents complete, the **Fund Manager** card will say **"awaiting you"**
5. Review the **Trader Decision panel** on the right — it shows:
    - Action (BUY / SELL / HOLD)
    - Quantity, Price Target, Stop Loss
    - AI Confidence % (needs > 65% to trigger a BUY)
    - Full Rationale written by the AI
6. **Approve, Reject, or Re-Analyze**

> **What happens when you click BUY NOW?**
> The Alpaca Executor fires a **Market GTC order** to Alpaca Paper Trading. You will receive an email confirmation immediately.

The Agent Log Stream

On the left side, the Agent Log Stream shows every step of the pipeline in real-time:

- `[TECHNICAL]` — The narrative interpretation of chart indicators
- `[NEWS]` — Impact rating and top headlines found
- `[SENTIMENT]` — Crowd psychology score from 4 data sources
- `[BULL]` / `[BEAR]` — Actual debate thesis text
- `[TRADER]` — The final structured trade decision
- `[RISK]` — Approval status and any adjustments made
- `[CORR]` — Correlation Guard check against your existing positions

---

4. The 9 AI Agents — Deep Dive

The pipeline runs in **3 phases**. Phase 1 agents run **in parallel simultaneously**.

PHASE 1 — Research (All 4 run at the same time)

🔴 Technical Analyst

Data source: Yahoo Finance (yfinance) — live OHLCV price data

What it computes:

- RSI (14-period)
- MACD + Signal Line + Histogram
- Bollinger Bands (upper, middle, lower)
- EMA 20 & EMA 50
- ATR (Average True Range — measures volatility)
- Trend direction (uptrend / downtrend / sideways)
- Support & Resistance levels

What the LLM does: Takes all the raw numbers and writes a precise 2-sentence interpretation. Example: *"With RSI at 64 and a bullish MACD crossover, momentum is turning upward; however, price sits just below EMA50 resistance which could cap near-term gains."*

AI model used: Groq gpt-oss-120b (Tier B — fast, ~1 second)

🔵 Fundamental Analyst

Data source 1: Yahoo Finance — P/E ratio, EPS, revenue growth, profit margins, debt/equity, analyst consensus

Data source 2: US Government SEC EDGAR — **Actual filing text**, not summaries:

- **10-K Annual Report** — Risk Factors section (what the CEO says could hurt the company)
- **MD&A** — Management Discussion & Analysis (CEO explains the financials in plain English)
- **8-K Emergency Filings** — Material events (CEO changes, lawsuits, product recalls, major contracts)
- **Form 4 Insider Trades** — When executives buy or sell their own company's stock

> ⚠️ **Note:** SEC EDGAR data is only available for US-listed stocks. For Indian stocks, crypto, and forex, the agent relies on Yahoo Finance metrics only.

Output: A health score (0-100) and valuation signal (undervalued / fairly valued / overvalued)

🟡 News Analyst

Data sources (no API key needed):

- Yahoo Finance RSS feed
- Seeking Alpha RSS feed
- Google Finance RSS feed
- Alpha Vantage (Fallback API redundancy)

Optional (if `NEWS_API_KEY` is set in `.env`):

- NewsAPI — adds premium sources (Reuters, Bloomberg headlines, etc.)

What it does: Filters all headlines down to those relevant to the specific ticker, then classifies the collective news impact as:

- `BULLISH` — positive news flow (earnings beat, new product, contract win)
- `BEARISH` — negative news (lawsuit, earnings miss, guidance cut)
- `NEUTRAL` — no material news

Time window: Last 3 days of news only. Old news is ignored.

🟢 Sentiment Analyst — Quad-Core Engine

The most sophisticated Phase 1 agent. It aggregates **4 independent data streams simultaneously**:

Stream	Source	What it reads
1. Finnhub	`finnhub.io` (API key needed)	Pre-calculated bullish/bearish article ratio, company news score, sector comparison
2. Alpaca/Benzinga	Alpaca News API (uses your existing keys)	Institutional-grade Benzinga headlines, sentiment tags

---

Stream	Source	What it reads
3. StockTwits	No API key needed	Retail crowd message volume + their pre-calculated sentiment (bullish/bearish tag ratio)
4. Google RSS	No API key needed	Mainstream media articles — fallback when other sources have limited data

> 🛡️ **FinBERT Noise Filter:** Before sending these 4 streams to the LLM, the system runs local HuggingFace `ProsusAI/finbert` to filter out pure noise (e.g. "TO THE MOON 🚀"). This protects the LLM from hallucinating based on retail hysteria.

Output: A unified "Crowd Psychology Score" synthesised by an LLM from all 4 simultaneously. This cross-source approach makes it resistant to single-source noise or fake news.

PHASE 2 — Adversarial Debate (Both run in parallel, then alternate)

🟢 Bull Researcher

AI model: Cerebras qwen-3-235b (deep reasoning, not speed-optimised)

Receives all 4 Phase 1 reports and builds the **strongest possible BUY thesis**. Rules it follows:

- Must cite specific metrics (not "the stock looks good")
- Writes 4-6 bullet points with actual data
- Ends with a specific price target rationale
- In debate rounds 2+: counter-argues the Bear's previous points

---

🔴 Bear Researcher

AI model: Cerebras qwen-3-235b

Takes the identical data and builds the **strongest possible SELL/avoid case**:

- Challenges every bullish assumption
- Points out what the Technical Analyst missed
- Escalates risks the Fundamental Analyst flagged
- The debate runs for 1-3 rounds, depending on conviction gap

> 💡 **Why debate?** A single AI model tends to be sycophantic (agrees with whatever data suggests). Two adversarial models forced to argue opposite positions surface risks that a single model would overlook.

LLM Model Assignment Table

TradeSense uses a **3-tier cascade system** — each agent tries models in priority order and auto-falls back if rate-limited. Gemini is always the absolute last resort (position #7) for every agent.

Agent	Tier	Primary Model	Fallback Chain
Fundamental Analyst	A — Strategic	Cerebras qwen-3-235b	→ Groq 120B → Groq 32B → Groq 20B → Llama4 → Cerebras 8B → Gemini
Bull Researcher	A — Strategic	Cerebras qwen-3-235b	→ Groq 120B → Groq 32B → Groq 20B → Llama4 → Cerebras 8B → Gemini
Bear Researcher	A — Strategic	Cerebras qwen-3-235b	→ Groq 120B → Groq 32B → Groq 20B → Llama4 → Cerebras 8B → Gemini
Trader Agent	A — Strategic	Cerebras qwen-3-235b	→ Groq 120B → Groq 32B → Groq 20B → Llama4 → Cerebras 8B → Gemini
Risk Manager	A — Strategic	Cerebras qwen-3-235b	→ Groq 120B → Groq 32B → Groq 20B → Llama4 → Cerebras 8B → Gemini
Technical Analyst	B — Analytical	Groq gpt-oss-120B	→ Groq 32B → Groq 20B → Cerebras 235B → Llama4 → Cerebras 8B → Gemini
News Analyst	B — Analytical	Groq gpt-oss-120B	→ Groq 32B → Groq 20B → Cerebras 235B → Llama4 → Cerebras 8B → Gemini
Fund Manager	B — Analytical	Groq gpt-oss-120B	→ Groq 32B → Groq 20B → Cerebras 235B → Llama4 → Cerebras 8B → Gemini
Scanner	B — Analytical	Groq gpt-oss-120B	→ Groq 32B → Groq 20B → Cerebras 235B → Llama4 → Cerebras 8B → Gemini
Sentiment Analyst	C — Speed-first	Groq gpt-oss-20B	→ Groq 120B → Groq 32B → Cerebras 235B → Cerebras 8B → Gemini

> **Why 3 tiers?** Tier A agents make decisions that move real money — wrong answer = wrong trade. They get the biggest, smartest model first. Tier C (Sentiment) does binary positive/negative classification — overkill to use a 235B model for that, so the fastest 20B model goes first.

🟣 Trader Agent

AI model: Cerebras qwen-3-235b (235B parameters — Tier A, highest priority model)

> Gemini is **NOT** used by the Trader Agent. It is position #7 — the absolute last-resort fallback used only if all Cerebras and all Groq keys are simultaneously rate-limited.

Reads **everything** — all 4 Phase 1 reports + the full Bull/Bear debate transcript — and synthesises a single structured decision:

```
{
  action:       BUY | SELL | HOLD
  quantity:     1-20 shares (paper trading limit)
  price_target: $XXX.XX  (the AI's future price estimate, NOT the limit order price)
  stop_loss:    $XXX.XX  (automatic exit if price falls here)
  confidence:   0.0 - 1.0
  time_horizon: intraday | swing | position
  rationale:    "2-3 sentence explanation of the decision"
  key_risks:    ["risk 1", "risk 2", "risk 3"]
}
```

Hard rule: The Trader will only emit BUY if `confidence > 0.65`. Below that, it emits HOLD.

🟠 Risk Manager

No LLM for the math checks — purely deterministic rules first, then LLM review.

Hard rules (automatic BLOCK if triggered):

- Confidence < 50% → BLOCK
- Quantity > 25 shares → REDUCE to 20
- BUY signal but trend is downtrend AND RSI > 75 → BLOCK (severely overbought in downtrend)

Soft rules (LLM judgment):

- Confidence 50–60% → Reduce position size by 30%
- RSI > 70 on a BUY → Flag elevated risk, allow only if fundamentals are strong
- Very high ATR (volatility) → Reduce position size by up to 50%

Correlation Guard (runs before LLM):

- Checks your existing open positions
- If the new trade correlates > 70% with something you already hold (e.g., AAPL and MSFT), it blocks or reduces to prevent concentrated exposure

---

⬜ Fund Manager — The Human Gate

This is where YOU step in.

In **HITL mode** (default):

- The LangGraph pipeline **pauses completely** here
- The dashboard shows you the full trade proposal
- You have 3 choices:
  - ✅ **BUY NOW / APPROVE** — Sends to Alpaca Executor
  - ❌ **REJECT** — Trade is cancelled, no order placed
  - 🔄 **RE-ANALYZE** — Starts the entire pipeline again from scratch

In **Autopilot mode** (when Autopilot is running):

- The LLM reviews the full proposal and auto-approves if `risk_score ≥ 0.55`
- You still receive an email notification immediately on any auto-approval

---

🟢 Alpaca Executor

The final node. Routes the order to the correct execution venue based on asset type:

Asset Type	Execution Venue	Order Type
US Stocks (AAPL, MSFT, NVDA...)	Alpaca Paper Trading	Market GTC
ETFs (SPY, QQQ, GLD...)	Alpaca Paper Trading	Market GTC
Crypto (BTC, ETH, SOL...)	Alpaca Paper Trading	Market IOC (fills instantly)
Forex (EUR/USD, USD/INR...)	In-house Forex Simulator	Simulated at live yfinance rates
Indian NSE Stocks (RELIANCE, TCS...)	NSE Simulator	Simulated at live yfinance rates

On every successful order: **sends you a confirmation email** with confidence score and AI rationale.

---

5. Debate Tab

After running any analysis, go to the **Debate** tab to read the full Bull vs Bear transcript.

This is especially useful if the AI said HOLD or REJECT — you can read exactly what the Bear researcher argued and decide if you agree.

The debate transcript is stored for your last analysis session. Running a new analysis overwrites it.

---

6. Positions & History

Positions Tab

Shows all currently open paper trades with:

- Entry price, current price, unrealised P&L
- Stop loss level
- Time held

The **Exit Manager** runs in the background monitoring these positions. It will automatically close a position and send you an email if:

- Price drops to the stop loss level → **Stop Loss exit (-5%)**
- Price reaches the AI's price target → **Take Profit exit (+12%)**
- Price hits a high and retraces → **Trailing Stop (+7% activation, +2% floor)**

History Tab

All closed trades with:

- Entry + exit price
- Final P&L (USD and %)
- The AI's rationale at time of entry
- Why it was closed (manual / stop loss / take profit)

---

7. Autopilot Mode

Autopilot runs the **full 9-agent pipeline autonomously** across 60+ assets every 30 minutes without you needing to be at your computer.

How it works:

1. The **AI Scanner** runs first — scores all 60+ assets from 0–100
2. Any asset scoring **≥ 60** becomes a "candidate"
3. For each candidate, the **full 9-agent pipeline** runs automatically
4. If the pipeline produces a BUY with confidence > 65% and risk approval, the Fund Manager **auto-approves** trades with risk_score ≥ 0.55
5. You receive an **Autopilot Cycle Summary email** at the end of every scan

To enable Autopilot:

Navigate to the **Autopilot** tab and toggle it on.

> ⚠️ Autopilot only trades Paper money. It requires valid Alpaca API keys to be configured in Settings.

---

8. AI Scanner

Click the **"AI Scanner"** button in the top nav to open the full scanner panel.

The scanner evaluates all 60+ supported assets and gives each a **score from 0–100** based on:

- Technical strength (trend, momentum, RSI health)
- News flow impact
- Sentiment signal

Assets scoring **≥ 60** are flagged as actionable.

The scanner processes assets in **batches of 20** internally to ensure the LLM output is never truncated and scores are always accurate.

---

9. Brain Tab — Knowledge Base

The Brain tab is your persistent intelligence layer — data that carries over between sessions and makes the AI smarter over time.

What it stores:

Section	What's in it	How it fills up
SEC Filings	10-K, 10-Q, 8-K, Form 4 text chunks from EDGAR	Fetch manually via API or it fills automatically as you trade US stocks
Learned Lessons	Post-mortem analysis written by AI after each closed trade	Fills automatically after every stop-loss or take-profit exit
Strategy Docs	PDFs or documents you upload	Upload via the Settings page
WS Events	Live autopilot events for the current session	Fills when Autopilot is running

Macro Environment Panel

Shows live macro data fetched from yFinance:

- 10Y & 2Y Treasury yields
- VIX (fear index)
- Fed Funds Rate
- Yield Curve spread
- CPI YoY

The **Bull/Bear researchers** use this macro context when building their theses.

To manually fetch SEC filings for a US stock:

```
POST https://tradesensebackend.skds.site/kb/fetch-sec/AAPL
```

Or locally: `curl -X POST http://localhost:8000/kb/fetch-sec/AAPL`

Once fetched, the **Fundamental Analyst** will automatically use this data in future analyses of that ticker.

---

10. Email Notifications

TradeSense sends **3 types of automated emails** to keep you informed without needing to watch the dashboard.

⚙️ How to Setup Email Notifications:

To receive emails, you must configure a free **Resend** API key via the dashboard settings.

1. Go to resend.com and create a free account (allows 3000 free emails/month).
2. Go to **API Keys** in Resend and click **Create API Key**.
3. Copy the token (it starts with `re_...`).
4. In TradeSense, open **Settings (⚙️ Gear Icon)**.
5. Expand the **Email Configuration** section.
6. Enter your `RESEND_API_KEY` and the `Receiver Email` (the email you want notifications sent to).
7. Click **Update Keys**. You can click **Test Email** to verify it is working.

📧 Email 1 — Trade Executed (fires immediately on every order)

```
Subject: TradeSense 🟢 BUY | AAPL | 📄 PAPER

Action:       🟢 BUY
Asset:        AAPL (stock)
Quantity:     5.0 shares
```

```
Est. Price:    $270.77 (~$1,354 total)
Confidence:    72%
Order ID:      [Alpaca order ID]
Mode:          📄 PAPER
AI Rationale: Strong technical momentum with bullish MACD crossover...
```

📧 Email 2 — Position Closed (fires on every stop-loss or take-profit)

```
Subject: TradeSense 💰 PROFIT | AAPL Closed

Ticker:        AAPL
Exit Trigger: Trailing Stop
P&L:           +$45.20 (+2.5%)
Quantity:      5
Entry Price:   $270.77
Exit Price:    $279.80
```

📧 Email 3 — Autopilot Cycle Summary (fires at end of every 30-min scan)

```
Subject: TradeSense 🤖 Autopilot Cycle | 2 trade(s) made

Assets Scanned:  60
Candidates:      4
Trades Executed: 2
Top Picks:       AAPL - 81/100 (BUY), NVDA - 77/100 (BUY)...
Next Scan In:    ~30 minutes
```

---

11. Supported Assets

US Stocks & ETFs (execute on Alpaca Paper)

`AAPL` `MSFT` `NVDA` `AMD` `TSLA` `META` `AMZN` `GOOGL` `PLTR` `SOFI` `SPY` `QQQ`

Crypto (execute on Alpaca Paper — fractional)

`BTC` `ETH` `SOL` `DOGE` `ADA` `AVAX` `XRP` `LINK` `LTC` `DOT`

Commodities & Global ETFs (execute on Alpaca Paper)

`GLD` (Gold)· `SLV` (Silver)· `USO` (Oil)· `UNG` (Natural Gas)· `GDX` (Gold Miners)· `COPX` (Copper)· `INDY` (Nifty 50 ETF)· `EWZ` (Brazil)· `FXI` (China)

Indian NSE Stocks (simulated via NSE Simulator using live yfinance rates)

`RELIANCE` `TCS` `INFY` `HDFCBANK` `ICICIBANK` `WIPRO` `BAJFINANCE` `TATAMOTORS` `ADANIENT` `HCLTECH`

- `NIFTY50` and `SENSEX` (analysis only, no execution)

Forex (simulated via in-house Forex Simulator at live yfinance rates)

INR pairs: `USDINR` `EURINR` `GBPINR` `JPYINR`

Global pairs: `EURUSD` `GBPUSD` `USDJPY` `AUDUSD` `USDCAD` `USDCHF` `NZDUSD`

---

12. Risk Management Rules

The Risk Manager enforces these rules **automatically** on every trade:

Rule	Threshold	Action
Low confidence	< 50%	BLOCK trade entirely
Position too large	> 25 shares	REDUCE to 20 shares
Overbought in downtrend	BUY + downtrend + RSI > 75	BLOCK trade entirely
Moderate confidence	50–60%	REDUCE quantity by 30%
High volatility	Very high ATR	REDUCE quantity up to 50%
High RSI on BUY	RSI > 70	Flag warning, allow if fundamentals strong
Correlated position	> 70% correlation with existing holding	Block or reduce

🛑 Exit & Protection Rules (Automatic)

Rule	Threshold	Action
Hard Stop Loss	-5.0%	Immediate Market Sell

Rule	Threshold	Action
Take Profit	+12.0%	Lock in gains
Trailing Stop (Level 1)	+7.0%	Activates trailing protection
Trailing Stop (Level 2)	Drop to +2.0%	Sell if price falls back to +2% after hitting +7%
AI Vital Sign Check	Every 30m	Risk Manager reviews if conviction has fundamentally collapsed

---

13. Tips & Best Practices

Getting the best results:

- Use the **Debate tab** after every analysis — the Bear researcher often flags real risks worth considering
- Check **Brain → SEC Filings** before trading a US stock. Fetch the filings if count is 0
- The **AI Confidence** threshold is 65%. If the AI says 62% and HOLD, trust it — that's exactly what the risk rules are for
- **Macro environment** matters — if VIX is above 25 (fear spike), the Bear researcher will be extra aggressive. That's intentional

Understanding HOLD decisions:

- HOLD means either: signals are genuinely mixed, OR confidence is below 65%, OR risk rules blocked it
- Click **Re-Analyze** if you want the AI to try again (sometimes different news data loads)

Running Autopilot:

- Start it during Indian market hours (09:15–15:30 IST) for NSE stocks
- Start it during US market hours (14:30–21:00 IST) for US stocks
- The 30-minute scanner cycle means it will catch intraday momentum shifts

Understanding emails:

- You get an **email for every single trade** — even if you are asleep
- Confidence 0% in an old email = that was a pre-fix issue. All new trades show the real % now

---

TradeSense — Built by Saswata Kumar Dash
Stack: LangGraph · Cerebras · Gemini · Groq · Alpaca Paper Trading · yfinance · SEC EDGAR · ChromaDB

---

14. Backtest Engine ⚗️

The Backtest Engine is TradeSense's **quantitative validation layer**. It replays the exact same mathematical signals that the live Autopilot uses against years of real historical price data — giving you statistical proof of whether the strategy works before touching real money.

How It Works

The engine runs a **walk-forward simulation** — it moves through history one trading day at a time, using only data that was available *on that day* (no look-ahead bias). At each day it:

1. Computes RSI, MACD, EMA20/50, Bollinger Bands and Volume scoring — the exact same Phase 1 formula the live scanner uses
2. If the score ≥ your threshold, it buys at the *next* day's open (realistic fill)
3. Monitors the position daily for stop-loss, take-profit, or trailing stop triggers
4. After the full simulation completes, sends the results to the AI for a single post-analysis report

> [!IMPORTANT]
> **No LLM calls during the simulation.** Only pure math — identical to live trading. This makes backtests free, instant, and 100% reproducible.

Controls

Control	What it does
Ticker	Any asset in the TradeSense registry (AAPL, BTC, TSLA, etc.)
Start / End Date	Your simulation period. Use at least 1 year for meaningful results
Initial Capital ($)	Virtual portfolio starting value
Position Size %	What % of available cash to deploy per trade (default 10%)
Stop Loss %	Hard exit if position drops this much. Mirrors live `exit_manager` rule
Take Profit %	Hard exit at this gain. Mirrors live rule
Min Score Threshold	Minimum Phase 1 scanner score (0–100) required to enter a trade

Performance Metrics Explained

Metric	What it means	Target
Total Return %	Portfolio growth over the full period	> 0% obviously

Metric	What it means	Target
Annualised Return %	Return expressed as a per-year figure	> 15% is excellent
Sharpe Ratio	Return per unit of risk (higher = better)	> 1.5 is excellent, > 1.0 is good
Sortino Ratio	Like Sharpe but only penalises *downside* volatility	> 1.0 is good
Max Drawdown %	Worst peak-to-trough drop	< 15% is safe
Win Rate %	% of trades that were profitable	> 55% is solid
Profit Factor	Gross profit ÷ gross loss	> 1.5 is good
Avg Hold Days	Average number of days a trade was held	Varies by strategy

The Equity Curve

The chart shows your virtual portfolio's value over time. A smooth, upward-sloping curve indicates a resilient strategy. Sharp drops indicate drawdown periods — study those periods (2020, 2022) to understand the strategy's weaknesses.

The AI Post-Analysis

After the simulation, a single LLM call generates a hedge-fund-style post-mortem covering:

- Overall performance quality versus professional benchmarks
- What market conditions the strategy excels in
- Where the strategy is vulnerable
- Specific, actionable optimization recommendations

Interpreting Results & What to Do Next

If Sharpe > 1.5 and Max Drawdown < 15%: The strategy is mathematically validated. You can have statistical confidence running the Autopilot with these parameters.

If Win Rate < 50% or Profit Factor < 1.0: The strategy is losing money historically. Try:

- Raising the min score threshold (be more selective)
- Tightening the stop-loss (cut losers faster)
- Testing a different date range (avoid testing only in bull markets)

Stress Testing — Recommended Date Ranges:

- `2020-01-01` to `2020-12-31` — COVID crash + V-shaped recovery
- `2022-01-01` to `2022-12-31` — Full bear market, tech down 70%
- `2019-01-01` to `2024-12-31` — Full cycle including bull, bear, recovery

> [!TIP]
> Run the same ticker with different **Stop Loss / Take Profit** combinations and compare Sharpe Ratios. The combination with the highest Sharpe is your optimal parameter set for that asset class.

> [!NOTE]
> The Backtest Engine only supports single-ticker simulations at this time. Multi-ticker portfolio backtesting is coming in a future release.

---

15. Switching to Real Money 💵

> [!CAUTION]
> **Real money means real losses.** Only switch after completing every checkpoint in the checklist below. There is no undo button on a live trade.

This chapter tells you exactly **when you are ready**, **what changes**, and **how to do it safely**.

When Are You Ready? — The 4-Stage Readiness Checklist

You must pass all four stages in order. Do not skip stages.

✅ Stage 1 — Paper Trading Stability (Weeks 1–4)

- [ ] Autopilot has been running on Paper for at least **3 full trading weeks** without crashes
- [ ] No hallucinated SELL events (reconciliation bug is confirmed fixed in your deployment)
- [ ] Crypto orders executing without `ORDER FAILED` errors
- [ ] You have received at least **5 real trade email notifications** confirming the pipeline is healthy
- [ ] Positions tab shows open positions that match what you see in Alpaca Paper dashboard

✅ Stage 2 — Backtesting Validation (Minimum 1 year of data)

Run the Backtest Engine on the assets you plan to trade live. You need ALL of these:

Metric	Minimum Threshold to Go Live
Sharpe Ratio	≥ 1.0
Max Drawdown	≤ 20%

Metric	Minimum Threshold to Go Live
Win Rate	≥ 52%
Profit Factor	≥ 1.2
Total Trades	≥ 15 (enough sample size)

> [!IMPORTANT]
> Run backtests on **at least 2 of these stress periods** before going live:
>
> - `2020-01-01` to `2020-12-31` (COVID crash)
> - `2022-01-01` to `2022-12-31` (full bear market)
>
> If the strategy survives these with Max Drawdown < 20%, it is fundamentally resilient.

✅ Stage 3 — Manual Confidence Check

- [ ] You understand exactly what RSI, MACD, Sharpe Ratio mean (not just trusting the AI blindly)
- [ ] You have read the AI Post-Analysis report from backtesting and understand the strategy's weaknesses
- [ ] You have set a **personal daily loss limit** (example: "If I lose more than ₹2,000 in a day, I stop")
- [ ] You have verified your Alpaca account has passed identity verification (required for live trading)

✅ Stage 4 — Micro Live Test

Before committing meaningful capital:

- [ ] Switch to live with `max_quantity = 1` (one share per trade maximum)
- [ ] Fund with a small amount you are 100% comfortable losing entirely (example: $50 / ₹4,000)
- [ ] Run for **2 weeks live** with this micro allocation. Confirm orders appear in Alpaca Live dashboard
- [ ] Confirm email alerts fire correctly for live orders
- [ ] Confirm exit_manager closes positions at stop-loss correctly using real fills

---

How to Switch — Technical Steps

Step 1 — Get Your Alpaca Live Keys

1. Go to Alpaca Live Dashboard
2. Complete **Identity Verification** — you need a government ID (Passport or Aadhaar)
3. Fund your account (minimum $1 to test, recommended $200–500 for initial live trading)
4. In the Alpaca dashboard → **API Keys** → Generate **Live** keys (they start with `PK`)

Step 2 — Update Settings in TradeSense

1. Open TradeSense → Click the **Gear Icon (⚙️)** → Settings
2. Paste your **Live API Key** and **Live Secret Key**
3. Click **Update Keys**

> [!WARNING]
> Your GCP VM's `.env` file also has `ALPACA_PAPER=true`. You must SSH into the VM and change this to `ALPACA_PAPER=false` for the backend to route orders to the live endpoint.

```
cd ~/backend
nano .env
# Change: ALPACA_PAPER=true → ALPACA_PAPER=false
sudo systemctl restart tradesense
```

Step 3 — What Changes Automatically

Component	Paper Mode	Live Mode
Alpaca endpoint	`paper-api.alpaca.markets`	`api.alpaca.markets`
Portfolio balance	Virtual $100,000	Your real USD balance
Order fills	Simulated	Real market execution
`exit_manager` sells	Test orders	Real market sell orders
Email alerts	Same	Same (now real trades)
Backtest Engine	Unchanged	Unchanged (always uses yfinance history, not broker)

Step 4 — Recommended Safety Config for Live

Make these changes **before** switching:

```
# backend/agents/risk_manager.py
MAX_QTY_USD = 50          # Never invest more than $50 per trade initially
MAX_POSITIONS = 3         # Reduce concurrent positions from 5 to 3

# backend/agents/exit_manager.py
STOP_LOSS_PCT = 4.0       # Slightly tighter stop when real money is at ris
```

Ongoing Risk Management (Live Trading)

- **Keep HITL ON** — Never disable Human-in-the-Loop for live trading. Every trade requires your click.
- **Monitor daily** — Check the Positions tab every morning. Real positions = real exposure.
- **Trust the stop-loss** — The Exit Manager will sell at -5% automatically. Do not override it emotionally.
- **Watch the Autopilot log** — If you see repeated `ORDER_FAILED` or `PIPELINE_ERROR` entries, stop the autopilot and investigate before it places more orders.
- **Scale slowly** — After 4 weeks of profitable live trading: increase `MAX_QTY_USD` by 50%. Never double overnight.

Realistic Timeline (Paper → Live)

Phase	Duration	Goal
Paper Trading	3–4 weeks	Confirm engine stability, zero crashes
Backtesting	1 week	Validate Sharpe ≥ 1.0, Drawdown ≤ 20%
Micro Live ($50)	2–3 weeks	Confirm live order flow, email alerts, exit manager
Small Live ($200–500)	4–6 weeks	Build confidence, refine stop-loss params from backtest
Normal Live	Ongoing	Scale position size based on proven backtested parameters

Estimated total time from today to confident live trading: 2–3 months.

> [!TIP]
> The safest way to scale: every month you are profitable, you are allowed to increase your capital allocation by 25%. Every month you are at a loss, you reduce it by 25%. Never yolo the whole balance.

TradeSense — Built by Saswata Kumar Dash
Stack: LangGraph · Cerebras · Gemini · Groq · Alpaca Paper Trading · yfinance · SEC EDGAR · ChromaDB