

CONFIDENTIAL ENTERPRISE PROJECT

AI Trading Assistant

Essential Product Management Documents - Portfolio Edition

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

Portfolio index

This master index packages a sanitized, generalized and presentation-ready product management case study for an AI trading assistant in a confidential blockchain mobile application.

Company	Confidential Blockchain Company
Role	AI Product Manager and Solution Architect
Project type	Confidential enterprise AI product
Confidentiality	High - NDA-protected / sanitized
Public evidence	Reconstructed visuals and generalized documents only

Document map

No.	Document	Primary content
00	Portfolio Overview	Narrative, evidence standard, document map and confidentiality model
01	Product Concept Document	Vision, problem, users, value, MVP scope, assumptions and risk
02	Product Requirements Document	Functional, AI, safety, data, analytics and pilot requirements
03	UX, Conversation and Service Design	Persona, journey, conversation patterns, mobile UX and usability
04	AI Strategy, Safety and Evaluation Plan	RAG/tools, financial safety, threat model, evaluation and monitoring
05	Architecture and Traceability Report	Layers, flows, trust boundaries, data handling and decisions
06	Pilot, Measurement and Roadmap Plan	KPIs, experiments, instrumentation, roadmap, backlog and risk register

Confidentiality notice

This document is a sanitized and generalized reconstruction created from memory for portfolio use. It excludes the company name, original product documents, source code, production prompts, customer or portfolio data, exact models, proprietary signals, trading logic, security controls, vendor details, performance metrics, and internal architecture.

Recommended public label

AI Trading Assistant - Confidential Blockchain Company - Sanitized Portfolio Case Study

Suggested short description

A polished, NDA-safe portfolio document set covering product strategy, investor journeys, conversational workflows, requirements, AI architecture, financial-safety guardrails, evaluation, roadmap, and human-controlled trade preparation for an AI-powered trading assistant in a blockchain mobile application.

Publishing guidance

- Publish the combined PDF and a separate thumbnail; keep editable Word files as interview backup.
- Label every diagram as conceptual, reconstructed or sanitized.
- Do not use the original company name, branding, real balances, strategies, prompts, API names or production metrics.
- Describe the product as decision support, not guaranteed financial advice or autonomous trading.

Interview talking points

- How user research and market-risk constraints shaped the MVP.
- Why the LLM was positioned as an explanation and orchestration layer rather than a source of truth.
- How market freshness, balanced evidence, confidence and explicit confirmation reduce user harm.
- How product, engineering, data, UX, compliance, legal and security share launch accountability.
- How the KPI framework avoids optimizing speculative trading activity as the primary outcome.

Do not disclose

- Original company documents, brand, internal team identities or unapproved dates.
- Real customer holdings, order history, transaction identifiers or conversation logs.
- Proprietary trading signals, ranking logic, prompts, models, providers, APIs or security controls.
- Exact production performance, adoption, revenue or trading-volume results.

CONFIDENTIAL ENTERPRISE PROJECT

AI Trading Assistant

Market intelligence and trading decision support for a confidential blockchain mobile platform

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

PORTFOLIO OVERVIEW

A portfolio-ready document set demonstrating AI product strategy, mobile UX, market-data orchestration, responsible financial guidance, architecture, evaluation, and delivery planning without exposing proprietary company information.

PRODUCT	EXPERIENCE	CONTROL
AI market intelligence assistant	Personalized research + risk-aware guidance	Human decision, explicit confirmation, auditability

Company	Confidential Blockchain Company
Role	AI Product Manager and Solution Architect
Product	In-app AI trading research and decision-support assistant
Primary channels	Mobile app, market dashboard, watchlist, portfolio views
Document status	Sanitized reconstruction from memory
Evidence standard	Actual contribution separated from illustrative recommendations

Confidentiality boundary

This document is a sanitized and generalized reconstruction created from memory for portfolio use. It excludes the company name, original product documents, source code, production prompts, customer or portfolio data, exact models, proprietary signals, trading logic, security controls, vendor details, performance metrics, and internal architecture.

Document map

Artifact	Purpose	Portfolio evidence
Portfolio Overview	Executive narrative and confidentiality framing	Product leadership and communication
Product Concept Document	Problem, target users, value, scope and business case	Discovery and strategy
Product Requirements Document	Functional, AI, data, safety and operational requirements	Requirements and traceability
UX, Conversation and Service Design	Journey, mobile UX, conversation flows and decision points	Customer experience design
AI Strategy, Safety and Evaluation	RAG, analytics, guardrails, evaluation and release gates	Responsible AI product design
Architecture and Traceability	Conceptual layers, data flows, controls and requirement mapping	Solution architecture
Pilot, Measurement and Roadmap	KPIs, experiments, rollout, backlog and risks	Delivery and product operations

Portfolio narrative

The case study presents a generalized AI assistant embedded in a blockchain trading application. It helps users understand market conditions, summarize relevant data, review portfolio exposure, explore technical and sentiment indicators, and prepare an order decision while keeping final judgment and execution under the user's control.



Sanitized mobile UI mockups reconstructed for portfolio use.

What this portfolio demonstrates

- Product framing across investor research, decision quality, trust, user education, and platform engagement.
- Mobile-first conversational UX linked to watchlists, portfolios, market data, and order-preview workflows.
- RAG, analytics orchestration, personalization, confidence controls, and explainable insight design.
- Financial-safety guardrails that prohibit guaranteed returns, autonomous execution, and unsuitable high-risk recommendations.
- Product requirements with testable acceptance criteria, observability, traceability, and rollout gates.

Product system at a glance

Layer	Representative capability	Product purpose
User context	Preferences, watchlist, portfolio, recent activity	Personalize research without assuming risk tolerance
Market intelligence	Prices, indicators, news, sentiment, on-chain signals	Create a current evidence set
AI reasoning	Intent detection, retrieval, synthesis, explanation	Turn data into understandable insights
Decision support	Scenarios, risk factors, order preview, alerts	Support informed human decisions
Controls	Freshness, disclaimers, consent, confirmation, audit logs	Reduce financial and operational risk

Confidentiality model

Label	Meaning
Reconstructed	Recreated from memory; not an original company artifact.
Illustrative	Shows a reasonable product pattern, not the exact implementation.
Target	A proposed goal or threshold, not a claimed production result.
Recommendation	A portfolio design choice added to explain best practice.
Verified contribution	A responsibility the author personally performed, expressed without sensitive detail.

Financial-use boundary

The assistant is positioned as market research and decision support. It does not promise returns, provide individualized regulated advice, or place trades without explicit user review and confirmation. Regional and product restrictions must be enforced by the platform.

Artifact usage guidance

- Publish the combined PDF as the primary portfolio artifact and keep editable files private unless requested.
- Label all images as sanitized, reconstructed, or conceptual.
- Do not add real account balances, transaction IDs, private strategies, internal prompts, exact models, or production metrics.
- During interviews, focus on trade-offs, decision logic, measurement, safety, and cross-functional leadership.

CONFIDENTIAL ENTERPRISE PROJECT

AI Trading Assistant

Product vision, investor problem, scope, value and business case

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

1. Executive summary

The AI Trading Assistant is a mobile decision-support product for retail cryptocurrency users. It consolidates current market data, portfolio context, technical indicators, news, sentiment, and educational explanations into an understandable conversational experience. The product is designed to improve research efficiency and decision quality while keeping execution, risk acceptance, and accountability with the user.

Core product promise

Help users understand what is happening, why it may matter, and what risks to consider - without predicting guaranteed outcomes or replacing the user's judgment.

2. Business problem

Problem	User impact	Business impact
Fragmented research across charts, news, social media and portfolio tools	Slow, inconsistent and emotionally driven decisions	Lower engagement and trust
Complex indicators and trading terminology	Users struggle to interpret signals	High education and support burden
Generic market commentary	Advice may not reflect watchlists or exposure	Low relevance and poor retention
Fast-moving, volatile markets	Stale information can mislead users	Reputational and compliance risk
Overconfident trading content online	Users may take excessive risk	User harm and regulatory scrutiny

3. Target users and jobs to be done

Persona	Primary job	Needs	Important boundary
Daniel - intermediate retail trader	Understand market conditions before making a decision	Current evidence, simple explanations, portfolio context	No guarantee or personalized investment promise
New investor	Learn concepts and evaluate an asset safely	Plain language, examples, risk education	Avoid leverage and advanced strategies by default
Active portfolio user	Monitor exposure, alerts and relevant developments	Watchlist, portfolio impact, rapid summaries	Do not infer suitability from holdings alone
Risk-conscious user	Understand downside and scenario risk	Volatility, concentration, liquidity, event risk	Present uncertainty and counter-evidence

User Journey Map

AI Trading Assistant

Persona

Name: Daniel

Role: Retail Cryptocurrency Trader

Experience: Intermediate Investor

Primary Goal: Make informed cryptocurrency investment decisions with the help of AI while maintaining full control over every trade.

Journey Map

Journey Stage	User Goal	User Actions	Thoughts	Emotion	Pain Points	AI Opportunities
1. Open AI Trading Assistant	Understand today's market	Opens the trading app and launches AI Assistant	"What's happening in the market today?"	🧐 Curious	Too much information from different sources	Personalized market overview and suggested prompts
2. Ask a Question	Receive investment guidance	Types questions like "Should I buy BTC today?" or "Analyze my portfolio."	"Can AI explain this in simple terms?"	👉 Interested	Doesn't understand technical indicators	Natural language understanding in conversations AI, personalized responses
3. AI Collects Context	Receive relevant insights	AI retrieves portfolio data, watchlists, recent trades, and real-time market information	"I hope it considers my portfolio."	🕒 Waiting	Generic recommendations are not useful	Context-aware analysis using portfolio, preferences, and market data

Reconstructed persona and early journey map.

4. Value proposition

UNDERSTAND	PERSONALIZE	CONTROL
Explain market movements and indicators in plain language	Connect insights to watchlists, preferences and portfolio exposure	Show uncertainty, risks, sources and explicit user confirmation

5. Product principles

Principle	Product interpretation
User decision first	The assistant informs; the user decides and confirms.
Evidence before narrative	Every material claim must be grounded in current approved data.
Risk before reward	Show downside, uncertainty and conflicting signals, not only upside.
Explain, do not obscure	Use understandable reasons, data timestamps and source types.
Personalize carefully	Use authorized context without inferring suitability or financial capacity.
Fail safely	Clarify, defer, or decline when data is missing, stale, restricted or unsafe.
Learning loop	Use feedback and outcomes to improve content, retrieval and UX - not to optimize risky behavior.

6. MVP scope

In scope	Out of scope for MVP
Conversational market summaries	Autonomous trading or portfolio rebalancing
Watchlist and portfolio-aware insights	Guaranteed forecasts, price targets or profit promises
Current price, volume, volatility and selected indicators	Unrestricted derivatives or leverage coaching
News and sentiment summaries with freshness indicators	Copy trading or social trading recommendations
Educational explanations and glossary support	Tax, legal or individualized regulated advice
Price/market alerts and user-configured monitoring	Use of private account data without explicit authorization
Order preview with risk reminders and explicit confirmation	Silent execution or background order placement

7. Primary use cases

Use case	Representative user question	Expected outcome
Market analysis	What is the outlook for BTC today?	Balanced summary with current data, risks, sources and uncertainty.
Asset comparison	Compare ETH and another asset for volatility and recent momentum.	Side-by-side evidence; no definitive buy/sell instruction.
Portfolio insight	How could today's market move affect my portfolio?	Exposure summary and scenarios without suitability claims.
Watchlist monitoring	Tell me when volume or volatility changes materially.	User-configured alert with explanation and timestamp.
Trade preparation	Show me a preview for this order.	Fees, slippage, exposure, risk warning and explicit confirmation.
Learning	Explain RSI, liquidity or liquidation in simple terms.	Clear education with examples and links to approved content.

8. Assumptions and validation plan

Assumption	Validation method	Evidence threshold
Users value one place for research and portfolio context	Prototype tests and structured interviews	Users complete research tasks with fewer context switches
Plain-language explanations improve confidence without increasing risky behavior	Usability test plus comprehension quiz	Higher comprehension; no increase in overconfidence indicators
Source transparency increases trust	A/B test cited vs uncited summaries	Higher trust and error detection without major latency penalty
Personalization improves relevance	Opt-in pilot with watchlist and portfolio context	Higher relevance rating with low privacy concern
Users accept guardrails and confirmation steps	Task tests for order preview and restricted requests	High completion and understanding of control boundaries

9. Proposed business case

Value driver	Product mechanism	Measurement approach
Engagement	Useful daily market and portfolio insights	Active use, return frequency, alert engagement
Trust	Grounded, transparent and risk-aware responses	Trust score, correction rate, complaint rate
Research efficiency	Reduced switching across tools	Task time and number of external context switches
Platform retention	Integrated research and execution preparation	Retention among opted-in users; not trading volume alone
Support reduction	Education and self-service explanations	Avoided support contacts for eligible questions

Responsible business metric

Trading volume or fee revenue must not be the sole success metric. Product health should balance user understanding, safe behavior, trust, accuracy, and long-term retention.

10. Product risks

Risk	Potential harm	Mitigation
Stale or incorrect data	Bad user decisions	Freshness checks, timestamps, source fallback, answer blocking
Overconfident language	User treats uncertainty as certainty	Calibrated language, scenarios, counter-signals and disclaimers
Personalization interpreted as advice	Suitability or regulatory concern	Separate context from recommendation; restrict individualized directives
Prompt injection or malicious content	Manipulated output or data leakage	Tool allowlists, content isolation, retrieval controls and red-team tests
Automation bias	User follows AI without review	Explicit user decision, confirmation and educational friction
Volatility events	Rapidly invalidated analysis	Real-time refresh, circuit breakers and incident mode

CONFIDENTIAL ENTERPRISE PROJECT

AI Trading Assistant

Functional, AI, data, safety and operational requirements

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

1. Product objective

Provide retail cryptocurrency users with a secure, current, explainable and risk-aware assistant that helps them research markets, understand portfolio exposure, monitor selected assets, and prepare decisions without replacing human judgment or executing trades autonomously.

2. Success definition

Dimension	Pilot definition	Evidence
User value	Users complete research tasks faster and understand the result	Task success, time, comprehension and relevance rating
AI quality	Responses are grounded, current, balanced and understandable	Evaluation set, human review and source audit
Safety	Unsafe financial advice and unauthorized actions remain below release tolerance	Adversarial testing and incident monitoring
Operational readiness	Latency, availability, fallback and observability meet pilot needs	Load tests, dashboards and runbooks
Trust	Users understand limitations and can inspect evidence	Trust survey, source engagement and correction rate

3. Epic and feature scope

Epic	MVP features	Later considerations
Conversational research	Intent detection, market summary, cited insights, clarifying questions	Multi-step research workspace
Market intelligence	Prices, volume, volatility, technical indicators, approved news/sentiment	On-chain analytics and advanced scenario tools
Personal context	Opt-in watchlist, preferences, portfolio exposure summary	Goal-based planning subject to regulatory review
Alerts	User-defined price, volatility and event alerts	Adaptive thresholds with explicit controls
Trade preparation	Order preview, fees, slippage, exposure and confirmation	Advanced order education; no autonomous execution
Learning	Definitions, explainers, examples and linked articles	Personalized learning paths
Safety and controls	Freshness, refusal, regional policy, audit logs, escalation	Expanded jurisdiction and product-policy engine

4. Functional requirements

ID	Requirement	Priority
FR-01	The user can ask a market or portfolio question in natural language.	Must
FR-02	The assistant detects intent and requests	Must

ID	Requirement	Priority
	clarification when necessary.	
FR-03	The system retrieves current authorized market data and displays its timestamp.	Must
FR-04	The assistant may use opt-in watchlist and portfolio context.	Must
FR-05	The response distinguishes facts, interpretation, uncertainty and risk.	Must
FR-06	The user can inspect source categories and open supporting information.	Must
FR-07	The user can configure and manage alerts.	Should
FR-08	The user can preview a trade with fees, estimated slippage and exposure impact.	Should
FR-09	Any order action requires explicit review and confirmation in the platform order flow.	Must
FR-10	The user can rate, report or correct an answer.	Must

5. AI behavior requirements

ID	Behavior	Acceptance rule
AI-01	Ground material market claims in retrieved data or approved content.	No unsupported factual claims.
AI-02	Use current timestamps and block analysis when required data exceeds freshness tolerance.	Freshness status visible.
AI-03	Present balanced evidence, including key risks and conflicting indicators.	No one-sided promotional narrative.
AI-04	Avoid guaranteed returns, certainty language and unqualified buy/sell directives.	Unsafe requests are redirected to education and risk framing.
AI-05	Do not infer user risk tolerance, wealth, suitability or financial need from account activity.	Personalization stays descriptive.
AI-06	Use concise plain language and define technical terms on request.	Comprehension appropriate for selected user level.
AI-07	State when the system cannot answer reliably and provide a safe fallback.	Deferral is preferred to guessing.
AI-08	Do not reveal internal prompts, hidden policies, credentials or private account data.	Security and privacy checks pass.

6. User stories and acceptance criteria

User story	Acceptance criteria
As an investor, I want a daily market summary so I can quickly understand major changes.	Uses current data; identifies drivers and counter-signals; shows timestamp; avoids guaranteed outlook.
As a portfolio user, I want to understand exposure so I can consider concentration and volatility.	Uses authorized holdings; explains concentration and scenario impact; does not label the portfolio suitable.
As a user preparing a trade, I want an order preview so I can understand costs and risks.	Shows asset, side, size, fees, slippage estimate, exposure change, warnings and confirmation step.
As a new user, I want technical terms explained simply.	Uses plain language, an example and a link to approved educational content.
As a safety reviewer, I want dangerous requests handled consistently.	Leverage, guaranteed profit, manipulation and evasion prompts trigger defined safe behavior.

7. Non-functional requirements

Category	Requirement
Availability	Pilot target and fallback behavior documented; critical market incident mode supported.
Latency	Simple queries meet an agreed percentile target; complex analysis shows progress rather than appearing frozen.
Security	Authenticated sessions, least-privilege tools, secret isolation, audit logs and abuse controls.
Privacy	Data minimization, explicit authorization, purpose limitation and regional retention rules.
Accessibility	Readable contrast, scalable text, screen-reader labels and non-color-only indicators.
Observability	Trace ID, tool calls, source timestamps, policy decisions, latency and outcome events captured.
Localization	Locale-aware numbers, time, language and jurisdictional disclaimers.

8. Financial safety and escalation policy

Trigger	Required behavior
Guaranteed profit, certainty or price-target demand	Refuse certainty; explain uncertainty and provide neutral evidence.
High-risk leverage or liquidation request	Provide educational risk warning; restrict coaching based on policy and region.
Suspected market manipulation or insider information	Decline assistance and surface policy guidance.
Data is stale, unavailable or conflicting	Do not generate a definitive conclusion; show limitation and retry options.
Unauthorized portfolio or account request	Do not retrieve data; request authentication or permission.
User distress or compulsive trading language	Use supportive language, promote pause controls, and surface responsible-trading resources where available.

9. Analytics events

Event	Key properties	Purpose
assistant_opened	entry point, locale, authenticated state	Adoption and funnel
query_submitted	intent, asset class, complexity, no raw sensitive text in analytics	Demand analysis
data_retrieved	source type, freshness, success/failure	Grounding quality
response_rendered	latency, citations, confidence band, safety route	Quality and performance
source_opened	source type and position	Trust and explainability
alert_created	type, user threshold, asset	Feature usage
order_previewed	asset, side, risk warning shown; no unnecessary sensitive details	Decision support
answer_feedback	helpful, inaccurate, unsafe, reason	Improvement loop
trade_flow_entered	from assistant yes/no; tracked separately from outcome optimization	Journey analysis

10. Definition of done for pilot

- Approved intents, markets, jurisdictions, tools and data sources are explicitly allowlisted.
- Offline evaluation meets quality, freshness, citation and safety release gates.
- Adversarial tests cover manipulation, prompt injection, guaranteed returns, leverage and data exfiltration.
- Order actions remain outside the model and require explicit confirmation in the existing platform flow.
- Dashboards, alerts, runbooks, fallback content and incident ownership are ready.
- Legal, compliance, security, privacy, support, product and engineering approvals are recorded.
- Pilot users see clear limitations, data timestamps and a way to report problematic responses.

No performance claim

All metric thresholds in this portfolio are illustrative targets. They are not disclosed production results.

CONFIDENTIAL ENTERPRISE PROJECT

AI Trading Assistant

Mobile UX, conversation flows, investor journey and decision-support design

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

1. Experience objective

Create a mobile experience that turns fragmented market information into a clear, transparent and risk-aware research flow. The assistant should help users move from question to evidence to decision preparation without presenting itself as an oracle or making the decision for them.

2. Persona

Name	Daniel - representative persona
Role	Intermediate retail cryptocurrency investor
Primary goal	Make informed decisions while maintaining full control over every trade
Needs	Current data, understandable analysis, portfolio context, risk visibility
Pain points	Information overload, technical language, conflicting signals, emotional decisions
Trust requirement	Sources, timestamps, uncertainty and explicit control boundaries

User Journey Map

AI Trading Assistant

Persona

Name: Dunkel

Role: Retail Cryptocurrency Trader

Experience: Intermediate Investor

Primary Goal: Make informed cryptocurrency investment decisions with the help of AI while maintaining full control over every trade.

Journey Map

Journey Stage	User Goal	User Actions	Thoughts	Emotion	Pain Points	AI Opportunities
1. Open AI Trading Assistant	Understand today's market	Opens the trading app and launches AI Assistant	"What's happening in the market today?"	🧐 Curious	Too much information from different sources	Personalized market overview and suggested prompts
2. Ask a Question	Receive investment guidance	Types questions like "Should I buy BTC today?" or "Analyze my portfolio."	"Can AI explain this in simple terms?"	👉 Interested	Doesn't understand technical indicators	Natural language understanding in conversations AI, personalized responses
3. AI Collects Context	Receive relevant insights	AI retrieves portfolio data, watchlists, recent trades, and real-time market information	"I hope it considers my portfolio."	🕒 Waiting	Generic recommendations are not useful	Context-aware analysis using portfolio, preferences, and market data

1

Sanitized persona and initial journey reconstruction.

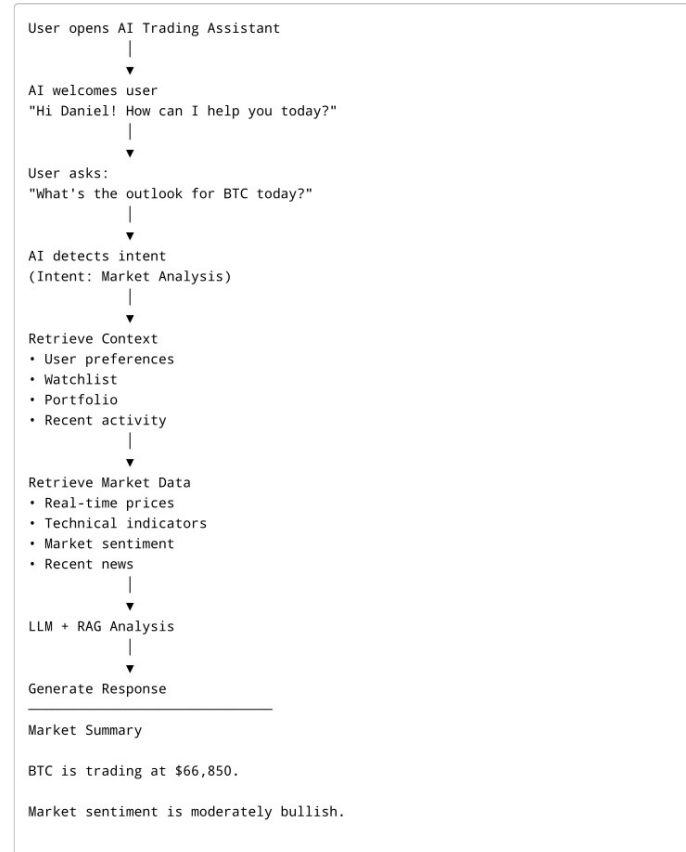
3. End-to-end user journey

Stage	User goal	User action	Pain point	Product opportunity
1. Open assistant	Understand today's market	Launch assistant from market, watchlist or portfolio	Too much information across tools	Context-aware entry and suggested prompts
2. Ask question	Receive relevant guidance	Ask in natural language	Does not understand technical indicators	Intent detection and plain-language clarification
3. Collect evidence	Understand what drives the answer	Review prices, indicators, news and portfolio context	Generic or stale commentary	Current, cited and personalized evidence set
4. Assess risk	Understand uncertainty and downside	Explore scenarios and counter-signals	Overconfidence and confirmation bias	Balanced risk cards and uncertainty language
5. Prepare decision	Review an order safely	Open order preview and review impact	Costs and exposure may be overlooked	Fees, slippage, exposure and confirmation
6. Monitor and learn	Track outcome and improve	Create alerts, save insight, provide feedback	No learning loop	Follow-up alerts, journal and education

4. Conversation workflow

AI Trading Assistant — Conversation Flow

Scenario 1: Market Analysis



1

Illustrative conversation flow for a market-analysis scenario.

Conversation step	UX behavior	Safety behavior
Welcome and context	Offer relevant prompts based on screen context without exposing private data.	Do not imply the system knows more than authorized context.
Intent and clarification	Identify asset, timeframe, question type and desired depth.	Clarify ambiguous terms before retrieving or summarizing.
Evidence collection	Show progress and data categories being checked.	Block stale or unauthorized sources.
Answer structure	Summary, evidence, counter-signals, risks, uncertainty, timestamp, next actions.	Avoid definitive prediction and guaranteed language.
Decision preparation	Allow watch, compare, set alert, learn more or open order preview.	No hidden or automatic order action.
Feedback and follow-up	Ask whether the answer was useful and allow correction/reporting.	Use feedback for quality, not to maximize risky trading.

5. Core mobile screens



Sanitized high-level mobile UI mockups.

Screen	Primary information	Design priority
Assistant home	Market snapshot, suggested prompts, alerts and watchlist context	Low cognitive load and clear limitations
Market insight	Summary, current data, sentiment, indicators, sources and risks	Evidence hierarchy and freshness
AI chat	Conversational analysis with chips and expandable details	Concise first, deeper evidence on demand
Portfolio insight	Balance, exposure, concentration, performance and scenarios	Descriptive, not suitability-based
Watchlist and alerts	Tracked assets, thresholds and alert reasons	User-controlled monitoring
Trade preview	Order details, fees, slippage, exposure change and confirmation	Friction proportional to risk

6. Conversation design patterns

Pattern	Example	Why it matters
Balanced answer	Momentum is positive, while volatility and event risk remain elevated.	Reduces one-sided framing.
Confidence language	The available signals are mixed; this is not a reliable short-term prediction.	Calibrates expectations.
Clarifying question	Do you want a short-term market summary or a longer-term educational overview?	Avoids answering the wrong question.
Evidence disclosure	Data updated at [time]; sources include market data, approved news and selected indicators.	Supports trust and error detection.
Risk-first next step	Review downside scenarios before opening the order preview.	Counters impulsive action.
Safe refusal	I cannot guarantee profit or provide hidden-insider information. I can explain current public signals and risks.	Maintains boundaries.

7. Decision-support service blueprint

Layer	Market analysis	Portfolio insight	Trade preparation
User action	Asks a question and reviews evidence	Opens exposure summary	Reviews order preview
Frontstage	Chat, charts, risk cards, sources	Allocation, concentration, scenarios	Fees, slippage, exposure impact, warning
Backstage	Intent routing and retrieval	Permissioned portfolio aggregation	Order service preview; no model execution
AI role	Summarize, explain and surface uncertainty	Describe exposure and possible scenarios	Explain fields and risks
Human control	Chooses what to trust and explore	Decides whether exposure is acceptable	Explicitly confirms or cancels
Evidence	Timestamps, source type, indicators	Authorized holdings and market data	Exchange quote and policy controls

8. Empty, error and edge states

State	User message	System behavior
No data	Current data is unavailable for this asset.	Do not infer; offer retry, another timeframe or educational content.
Stale data	The latest verified data is older than the allowed freshness window.	Block definitive analysis and show timestamp.
Unsupported asset	This asset is not supported in the current experience.	Avoid external speculation; offer general education.
Regional restriction	This feature is not available in your region.	Apply policy before model generation.
Tool failure	I could not retrieve one of the required data sources.	Show partial-data warning or safe fallback.
Unsafe request	I cannot help guarantee returns or evade platform controls.	Refuse and redirect to neutral education.

9. Usability test plan

Task	What to observe	Success indicator
Interpret a market summary	Can the user separate facts, interpretation and uncertainty?	Correct comprehension without overstated confidence
Find supporting evidence	Can the user locate timestamp and sources?	Evidence found quickly
Review portfolio exposure	Does the user understand concentration and scenario limitations?	Accurate explanation in own words
Prepare a trade	Does the user notice fees, slippage, exposure and confirmation?	No missed critical control
Handle a refused request	Does the user understand why and know the safe alternative?	Boundary accepted without confusion

10. Accessibility and inclusive design

- Use plain language by default and allow users to request more technical depth.
- Do not rely only on red/green or up/down colors to communicate risk.
- Support screen readers, scalable text, reduced motion and keyboard/switch navigation where applicable.
- Localize time, currency, decimals, terminology and risk disclosures.
- Avoid shaming language around losses, mistakes or inexperienced questions.

CONFIDENTIAL ENTERPRISE PROJECT

AI Trading Assistant

AI approach, financial-safety guardrails, model evaluation and responsible operation

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

1. AI product strategy

Use a language model as an orchestration and explanation layer over approved market, portfolio and educational data. The model should not be the source of truth for prices, balances, orders, policy decisions or execution. Deterministic services remain authoritative for all financial data and actions.


RETRIEVE	SYNTHESIZE	CONTROL
Current, authorized evidence	Explain and compare with uncertainty	Policy, confirmation and monitoring

2. AI capability map

Capability	Method	Primary control
Intent and entity understanding	Classifier/LLM with structured output	Schema validation and ambiguity threshold
Market-data retrieval	Approved tools and APIs	Allowlists, freshness and source health
Knowledge retrieval	RAG over approved educational and policy content	Versioning, access control and citations
Analysis synthesis	LLM prompt with structured evidence package	Balanced-answer template and risk section
Portfolio explanation	Permissioned aggregation plus deterministic analytics	Data minimization and no suitability inference
Alert explanation	Rules/analytics generate trigger; LLM explains	Reason code remains authoritative
Order support	Deterministic order preview; LLM explains fields	No order tool available to the model

3. Reference response pipeline

1. Classify intent, entities, timeframe, market and requested depth.
2. Apply jurisdiction, product, age/account and feature eligibility policies before retrieval.
3. Retrieve authorized market data, portfolio context and approved knowledge.
4. Validate freshness, completeness, conflicts, units and asset identity.
5. Build a structured evidence object with facts, indicators, risks and source metadata.
6. Generate a response using a fixed answer schema and calibrated language.
7. Run policy, grounding, sensitive-data and action checks.
8. Render sources, timestamp, limitations and user-controlled next steps.
9. Record trace and outcome events for monitoring and review.



Illustrative end-to-end product workflow and AI touchpoints.

4. Grounding and data management

Data class	Examples	Grounding rule
Authoritative account data	Portfolio, balances, order preview, preferences	Only from authenticated internal services; never generated by the model.
Authoritative market data	Price, volume, order book, volatility, indicators	Use approved providers with timestamps and health checks.
Curated knowledge	Product guides, risk education, policies, glossary	Versioned RAG corpus with ownership and review dates.
External information	Approved news and public sentiment feeds	Summarize with provenance and avoid treating sentiment as fact.
Derived analytics	Trend, concentration, scenario, technical indicator calculations	Deterministic calculation service where possible.

5. Tool-use policy

Tool category	Allowed use	Prohibited use
Market data	Read current approved market information	Fabricate or silently substitute stale data
Portfolio	Read authorized descriptive exposure data	Infer wealth, suitability or financial need
Watchlist and alerts	Create or modify after user confirmation	Create hidden alerts or manipulative nudges
Order preview	Retrieve a non-binding preview	Place, cancel or modify an order autonomously
Education and policy	Retrieve approved guidance	Override jurisdiction or product restrictions
External content	Read allowlisted sources	Browse arbitrary untrusted links or execute embedded instructions

6. Guardrails and threat model

Threat	Example	Control
Hallucinated market fact	Invented price or event	Authoritative tools, structured evidence, post-generation claim check
Stale analysis	Outdated data during volatility	Freshness TTL, live refresh and answer blocking
Prompt injection	News article instructs model to ignore policy	Content/data separation, instruction hierarchy and tool allowlists
Market manipulation assistance	Request to coordinate pump-and-dump	Policy refusal, abuse signal and escalation
Guaranteed-return claim	User asks for certain winning trade	Refusal plus uncertainty and education
Data leakage	Prompt asks for another user's portfolio	Authentication, authorization and privacy filters
Automation bias	UI makes answer look authoritative	Source hierarchy, uncertainty, user decision language and confirmation
Compulsive trading risk	User asks repeated high-risk rapid trades	Responsible-use friction, limits and support resources where applicable

7. Evaluation framework

Evaluation dimension	Illustrative measure	Review method
Grounded accuracy	Claim support and numerical consistency	Labeled test set plus deterministic checks
Freshness	Age of data used vs policy	Automated trace audit
Completeness	Required answer sections present	Schema validation
Balance	Risks and counter-signals represented	Human rubric
Calibration	Language matches evidence uncertainty	Human rating and challenge set
Financial safety	Unsafe advice, certainty, leverage and manipulation handling	Red-team suite
Privacy and security	Unauthorized data and prompt-injection resistance	Security testing
Usefulness	Task success, comprehension and relevance	User study and pilot feedback
Latency and reliability	Percentile latency, tool failure and fallback	Operational dashboards

8. Illustrative release gates

Gate	Required evidence	Decision owner
Data integrity	Freshness and numerical checks pass for supported markets	Data and engineering
Grounding	Material claims meet support threshold	AI product and evaluation
Safety	No critical failures in manipulation, certainty, leverage or privacy tests	Compliance, legal and security
UX comprehension	Users understand uncertainty, sources and control boundaries	Product and UX
Operational readiness	Monitoring, runbooks, fallback and incident ownership ready	Engineering and operations
Scope approval	Jurisdictions, markets, users and tools documented	Product, compliance and legal

9. Human review program

- Sample responses by intent, asset, language, risk level and model version.
- Require specialist review for financial-safety, regulatory and manipulation scenarios.
- Review both AI content and the evidence package that produced it.
- Track recurring failure modes to owners with remediation dates.
- Re-evaluate after model, prompt, data source, policy or product changes.

10. Monitoring and improvement loop

Signal	Action
Spike in stale-data blocks	Investigate provider health, freshness thresholds and fallback.
Rise in user corrections	Sample traces, classify errors and update retrieval/evaluation.
Unsafe-language alert	Disable affected path if needed and start incident review.
Asset or event-specific anomaly	Activate incident mode and tighten answer scope.
Model or prompt release	Run regression, canary and post-release sampling.
User feedback	Improve relevance and clarity without optimizing speculative trading frequency.

Safety ownership

The product manager owns the product policy and release trade-offs, but launch approval requires shared accountability across AI, engineering, compliance, legal, security, data, operations and UX.

Conceptual architecture, data flows, controls and requirement traceability

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

1. Architecture intent

The architecture separates authoritative financial systems from generative AI. The assistant orchestrates retrieval and explains structured evidence, while market data, portfolio state, calculations, policy decisions and order actions remain under deterministic services and existing platform controls.

2. Sanitized high-level architecture



Conceptual architecture reconstructed from memory. No proprietary system details are shown.

Architecture disclaimer

Component names, boundaries and technologies are generalized. The diagram communicates product responsibilities, not the company's exact production design.

3. Layer responsibilities

Layer	Responsibilities	Key controls
Client layer	Mobile assistant, market context, charts, alerts, order preview links	Authenticated session, consent and secure rendering
Application layer	Gateway, intent routing, context, orchestration, recommendations and	Rate limits, schemas, workflow state and feature policy

Layer	Responsibilities	Key controls
	notifications	
AI and intelligence	LLM, RAG, analytics synthesis, personalization and explanation	Grounding, prompt controls, calibration and evaluation
Data layer	Market data, news, sentiment, portfolio, watchlists, historical features	Freshness, lineage, entitlement and data quality
Integration layer	Market providers, trading services, portfolio services, notifications and content	Allowlists, timeouts, retries and circuit breakers
Infrastructure and governance	Cloud, observability, security, model registry, feature flags and audit	Least privilege, change control, incident response and retention

4. Primary data flows

Flow	Sequence	Control points
Market question	User query -> intent -> policy -> market/news retrieval -> evidence validation -> LLM synthesis -> safety check -> response	Asset identity, timestamp, source health, claim support
Portfolio insight	Authenticated query -> permission -> portfolio aggregation -> exposure analytics -> synthesis -> response	Authorization, minimization, no suitability inference
Alert creation	User instruction -> parse threshold -> preview -> confirm -> alert service	Explicit confirmation and editable settings
Order preparation	Assistant insight -> order preview request -> deterministic quote/risk data -> user review -> existing order confirmation	Model has no execution credential or order command
Feedback	User rating/report -> trace link -> taxonomy -> review queue -> remediation	Sensitive text handling and reviewer access

5. Trust boundaries

Boundary	Risk	Architecture response
Mobile client to platform	Session theft, tampering	TLS, token validation, device/session controls
Orchestrator to AI provider/model	Sensitive context or prompt exposure	Minimize inputs, redact, isolate tenants, contractual controls
AI to tools	Unauthorized actions	Least-privilege read tools, schemas and policy middleware
External content to retrieval	Prompt injection and misinformation	Allowlist, sanitize, separate data from instructions
Portfolio services	Private financial data	Purpose-bound authorization and audit logs
Order system	Financial action risk	No direct model execution; existing transaction controls and user confirmation

6. Data classification and handling

Class	Examples	Handling
Public market data	Prices, public indicators, approved news	Cache with freshness and provenance.
Customer confidential	Portfolio, watchlist, preferences, conversation history	Encrypt, minimize, purpose-limit, retain by policy.
Restricted platform data	Risk controls, eligibility, internal policy and security signals	Never expose in model output; strict access and audit.
Model telemetry	Prompt template version, tool calls, latency, safety decision	Limit raw text; tokenize/redact sensitive fields.
Portfolio artifacts	Evaluation data and reconstructed diagrams	Synthetic or sanitized only.

7. Reliability and fallback design

Failure	Fallback
Market provider unavailable	Use approved secondary source or block current analysis with transparent message.
One indicator service fails	Provide partial-data answer only if policy allows and label the limitation.
LLM unavailable	Show deterministic market snapshot, educational content and manual navigation.
Policy service unavailable	Fail closed for personalized or action-related flows.
Portfolio service unavailable	Do not infer holdings; offer general market information.
High volatility incident	Reduce scope, tighten freshness, disable selected recommendations and use incident messaging.

8. Requirement traceability

Requirement	Architecture component	Verification
FR-03 current market data	Market data adapters + freshness validator	Timestamp and stale-data tests
FR-04 authorized context	Consent + portfolio service adapter	Permission and negative-access tests
FR-06 source transparency	Evidence store + UI source panel	Source rendering test
FR-09 explicit confirmation	Order preview integration only	No execution tool; end-to-end confirmation test
AI-03 balanced evidence	Evidence schema + response template	Human rubric and challenge set
AI-04 no guaranteed advice	Policy engine + safety classifier	Red-team tests
NFR security	Gateway, IAM, audit and secrets management	Security review and penetration test

9. Architecture decision records

Decision	Why	Trade-off
LLM is not a market-data source	Protect factual integrity	More orchestration and validation complexity
No autonomous order execution	Maintain user control and reduce financial risk	Additional user steps
Structured evidence before generation	Improve grounding and auditability	Higher latency and engineering effort
Policy before retrieval and generation	Block disallowed requests early	Requires robust policy service
Opt-in personalization	Respect privacy and autonomy	Less relevance for anonymous use

10. Limitations of this portfolio architecture

- No real vendor, model, database, API, cloud or security technology is identified.
- No proprietary strategy, signal construction, ranking logic or trading workflow is reproduced.
- No production capacity, latency, cost, accuracy or business results are claimed.
- The design intentionally emphasizes product boundaries and responsible controls over implementation specificity.

CONFIDENTIAL ENTERPRISE PROJECT

AI Trading Assistant

Metrics, experimentation, rollout gates, roadmap and product risk

SANITIZED	RECONSTRUCTED	PORTFOLIO USE
-----------	---------------	---------------

1. Measurement strategy

Measure whether the assistant improves user understanding and research efficiency without increasing unsafe behavior or overconfidence. Product success must balance user value, AI quality, safety, trust, operational health and long-term platform outcomes.

2. KPI tree

Outcome	Leading indicators	Guardrails
Better-informed decisions	Task success, comprehension, evidence use, research time	Overconfidence, unsafe advice and correction rate
Trusted assistant	Helpful rating, source engagement, repeat use	Complaint rate, misleading-language rate
Efficient research	Fewer context switches, faster insight completion	Latency and abandonment
Responsible engagement	Alert usefulness, educational use, voluntary return	Compulsive-use signals and speculative-volume optimization
Operational readiness	Availability, tool success, freshness and fallback	Incidents, data leakage and policy failures

3. Metric definitions

Metric	Definition	Interpretation caution
Research task success	User reaches a correct and understood answer for a defined task	Do not equate chat completion with understanding.
Grounded claim rate	Material claims supported by retrieved evidence	Source quality still requires review.
Freshness compliance	Responses using data within policy window	Different intents need different freshness.
Balanced-answer rate	Required risks/counter-signals present	Presence alone does not ensure usefulness.
Unsafe advice rate	Responses violating financial-safety policy	Critical failures require severity weighting.
Research time	Time from question to accepted evidence set	Faster is not better if comprehension falls.
Trust score	User confidence in reliability and boundaries	Avoid measuring blind trust as a positive outcome.
Order-flow entry	Users who open order preview after assistant use	Track for journey insight, not as the primary objective.

4. Pilot hypotheses

Hypothesis	Test	Pass signal
Structured summaries reduce research effort	Compare existing workflow with assistant prototype	Lower task time with equal or better comprehension
Citations improve trust and error detection	A/B source display patterns	Higher appropriate trust and issue reporting
Portfolio context improves relevance	Opt-in contextual vs generic response	Higher relevance without privacy concern
Risk-first design reduces overconfidence	Test balanced vs promotional-style	Better uncertainty calibration

Hypothesis	Test	Pass signal
	presentation	
Order preview supports informed control	Usability test with fees and exposure impact	Critical fields noticed before confirmation

5. Pilot design

Element	Illustrative plan
Audience	Small opted-in cohort in approved jurisdictions and supported spot markets.
Scope	Selected intents: market summary, asset explanation, portfolio exposure, alerts and order preview education.
Duration	Time-boxed pilot long enough to include normal and volatile market conditions.
Comparison	Baseline task workflow or staggered rollout where ethically and operationally appropriate.
Review	Daily operational checks, weekly quality review and immediate critical-incident escalation.
Exclusions	Unsupported assets, restricted products, autonomous execution and unapproved individualized advice.

6. Instrumentation plan

Funnel stage	Events	Questions answered
Discover	assistant opened, suggested prompt selected	Where and why do users enter?
Ask	query submitted, clarification requested	Which intents and ambiguities matter?
Retrieve	tool success, source freshness, conflicts	Is the evidence reliable and current?
Understand	response rendered, source opened, expand details	Do users inspect and understand evidence?
Act	alert saved, comparison opened, order preview opened	Which safe next steps are useful?
Reflect	feedback, correction, report, follow-up query	What failed and what should improve?

7. Phased roadmap

Phase	Product scope	Exit gate
0. Discovery	User research, workflow mapping, policy and data inventory	Approved opportunity and risk framing
1. Internal prototype	Synthetic data, limited intents, evaluator tools	Core grounding and safety pass
2. Employee alpha	Read-only current data, educational use, trace review	Operational monitoring and issue closure
3. Limited pilot	Opt-in users, selected assets and contextual insights	User value plus safety and reliability gates
4. Controlled expansion	More markets, languages, alerts and portfolio scenarios	Regional, capacity and evaluation readiness
5. Mature operation	Continuous evaluation, governance and incident learning	Sustained product health and change controls

8. Prioritized backlog

Priority	Item	Rationale
P0	Authoritative data tools, freshness checks and	Foundation for factual integrity

Priority	Item	Rationale
	response schema	
P0	Financial-safety policy, refusal and confirmation design	Required before user access
P0	Traceability, evaluation set, monitoring and incident controls	Required for enterprise operation
P1	Market summary, educational explanation and source panel	Core user value
P1	Opt-in watchlist and portfolio exposure context	Relevance and retention
P1	Alert creation and explanation	Ongoing monitoring use case
P2	Advanced comparisons, scenarios and multilingual depth	Expansion after core quality
P2	Trading journal and reflective learning tools	Long-term decision quality

9. Risk register

Risk	Likelihood	Impact	Treatment
Incorrect or stale financial information	Medium	High	Authoritative sources, freshness block, circuit breaker
Unsafe individualized advice	Medium	High	Policy, evaluation, scope limits and human governance
Prompt injection or data leakage	Medium	High	Security isolation, allowlists and red-team testing
Overconfidence and automation bias	Medium	High	Balanced design, uncertainty, sources and user confirmation
Low relevance or generic answers	Medium	Medium	Clarification, context, retrieval tuning and feedback
High latency or provider failure	Medium	Medium	Caching, fallbacks, progress UI and graceful degradation
Regulatory or regional mismatch	Low/Medium	High	Policy service, jurisdiction scope and legal review
Metric gaming toward trading activity	Medium	High	Balanced scorecard and governance review

10. Go / no-go checklist

- Supported markets, intents, jurisdictions and users are documented and enforced.
- Grounding, freshness, balanced-answer and safety gates pass on the frozen evaluation set.
- Critical red-team failures are zero and high-severity issues have approved treatment.
- Order execution remains outside the model and explicit user confirmation is verified.
- Dashboards, on-call ownership, runbooks, feature flags and fallback experiences are ready.
- Privacy, security, compliance, legal, data, UX, support and product approvals are recorded.
- Pilot consent, disclosures, feedback and complaint channels are available.

11. Portfolio outcome statement

Generalized outcome

The work demonstrates the ability to shape a complex AI-finance product from problem framing through UX, requirements, architecture, evaluation, safety and rollout planning. Specific production adoption, revenue, trading, accuracy and performance results are intentionally not disclosed.