Practical AI for Markets and Trading in Financial Services

Archive text edition of the revised manuscript · 11 November 2026

Prepared by the Harbour Think Tank under the leadership of Adrian Dunkley. The original evidence review is dated 26 October 2026.

About this paper

The Harbour Think Tank of the Caribbean AI Association covers the economy, work and enterprise. This paper applies that remit to markets and trading operations, one of fourteen functions in a financial-services collection. [1]

We wrote this paper for dealing-desk supervisors, execution analysts, trade-support teams, settlement officers and control specialists. Every proposal in it prepares evidence or investigates an operational question. The proposals exclude investment selection and trade signals, and they exclude order placement. Portfolio construction, issuer transactions, enterprise risk management and wider compliance programmes have their own papers.

This revision adds a page of build notes to each use case. The notes give a starter prompt, example tools, skills, a method and the limits for an AI agent.

Foreword

A price without a timestamp can be of no use to explain a trade. An execution message records that a trade occurred. It says nothing about whether the securities and the cash settled. Be careful with an attractive daily summary, because it can hide the one open exception that needs attention before a deadline.

An assistant can help to put these records in order while the source systems stay authoritative. The desk must be able to see what the tool read and where it did not reconcile. The desk must also see who accepted its explanation.

Start away from live order entry. Test one read-only work product against records whose answers you already know, with the difficult cases among them. Keep operational authority with the responsible people.

Adrian Dunkley, President, Caribbean AI Association

Acknowledgements and terms

The Harbour Think Tank of the Caribbean AI Association prepared this paper, and President Adrian Dunkley led the work. The paper gives ten proposed pilots, with prompts and agent levels. No one has run them yet. They are not trading strategies, and we claim no improved returns. The sources give context and operating disciplines, and a citation implies no endorsement by the publisher.

An artificial intelligence (AI) assistant finds and organises information, and it can draft text. An agent can also do configured steps with connected tools. The agents in this paper have no authority to enter or cancel orders, and no authority to make payments or to settle. Execution completes a trade in the market. Settlement discharges the delivery and payment obligations that follow. Reconciliation compares records to find differences. An exception is an item that needs investigation, and an exception accuses no one.

Market liquidity is the ability to transact without unacceptable delay or price impact. A benchmark is a defined reference point for comparison, and it is not always a price at which anyone could trade. In real operations, the technical and legal definitions in the applicable rules and contracts govern.

A prompt is the written instruction that a person gives to an assistant. The word skill has two meanings in this paper. A staff skill is an ability that a person learns. An agent skill is a saved set of instructions that an assistant loads for one task. [8]

Approved source

Draft or retrieval

Human review

Accepted output

The operating boundary

The assistant gives read-only evidence and internal drafts only. Every decision on investment, orders, allocation, payment, settlement and control stays with authorised people. Existing alerts and deadlines operate independently.

Contents

Section | Page

Executive summary | 5

1 The regional business case | 6

2 What the evidence supports | 7

3 The enterprise lens | 8

Process design: BPMN swimlanes | 9

4 Ten practical use cases: how to read each use case | 10

Use cases 01 to 10, each with build notes | 11-30

Priority and responsibility tables | 31-32

5 The foundations lens | 33

6 From pilot to practice | 34

Appendix A Pilot decision canvas | 35

Appendix B Measure value | 36

References | 37-38

Read the summary and the priority tables first. Then select one use case and define a pilot. Complete the decision canvas with the people who do the work.

Executive summary

Start with a source-linked pre-market brief and checks on a licensed market-data snapshot, then add a replay of a closed incident. Test them with no connection to trading instructions or settlement instructions.

Quantities and amounts stay with fixed calculations, and limit tests stay with existing control systems. Generated text can explain a checked result. It cannot change the calculation or hide a warning, and it cannot decide how the desk trades. Keep the time that data arrived apart from the time of the event, and keep the two apart from the processing time.

The priority table screens pilots under stated assumptions. Before the test, approve the input records, the reviewer, the quality threshold and the stop condition. Then count the missed exceptions, the unsupported explanations and the source errors, and total the review effort. Keep trading profit and favourable price moves out of the assessment. Neither shows that the assistant did the operational work better.

Each use case now has build notes with a starter prompt and example tools. The notes also set a limit for AI agents. In the pilots, no agent acts without the approval of a named person.

Four decisions for the sponsor

Select the desk work product. Name its owner. Set the boundary around the live systems and actions that the assistant can never reach. Agree what keeps alerts and required processing in operation on a day when the assistant fails or the desk rejects its output.

Choose a bounded process

Measure the baseline

Test with review

Decide from evidence

1 The regional business case

The Payment System Data Bulletin of the Bank of Jamaica (BOJ) for May 2025 reports key transaction volumes in the JamClear Central Securities Depository (JamClear-CSD). [2] These are historical counts. They do not measure investors or market prices, and they do not measure settlement failures.

Source: Bank of Jamaica, Payment System Data Bulletin, May 2025, Table CYSU05, page 6. The contents page gives the date January 2026. The population is the key activity categories that the table records: repurchase agreements (repos), pledge, trustee delivery-versus-payment, purchase and sale, and free-of-payment transfers. These are counts, and they are not a percentage with a population denominator. They do not cover all Jamaican securities trades, and they are not an estimate for 2026. [2]

No published infrastructure count shows one dealer how much reconciliation work it has, or how many errors an assistant can prevent there. Measure the records and exception categories of the firm, and leave market-size proxies out of the case.

The BOJ says that JamClear-CSD records the ownership of BOJ and Government of Jamaica domestic securities. Settlement goes through JamClear-RTGS, the real-time gross settlement system of the BOJ. In the description of the BOJ, settled RTGS payments are final and irrevocable. [3] Because of that finality, a drafted explanation of a payment status must stay apart from any authority to start or change a payment instruction.

Jamaica: JamClear-CSD key transaction volumes

Jan 2025: 4441

Feb 2025: 3749

Mar 2025: 4475

Apr 2025: 4320

May 2025: 4780

Bank of Jamaica, Payment System Data Bulletin, May 2025, Table CYSU05, page 6. The contents page has the date January 2026. [2] These are counts, so the chart shows no rate. The population is the key categories of the table: repurchase agreements (repos), pledge, trustee delivery-versus-payment, purchase and sale, and free-of-payment transfers.

2 What the evidence supports

The Main Market Rule Book of the Jamaica Stock Exchange (JSE) has amendments that connect contract notes and transaction clearing to the prevailing settlement cycle. The definition and the amendments that we inspected took effect on 27 May 2024. [4] For this reason, a desk must keep current operating instructions for each market. Do not let a model recall a settlement period, because its answer can be out of date. Confirm the applicable cycle and calendar, with any exceptions, before live use.

The Market Conduct Guidelines of the Financial Services Commission (FSC) of Jamaica cover client interests, fair dealing, conflicts, confidentiality and the line between fact and opinion. [5] They give a reviewer reason to read execution explanations and communications with care. Nothing in them suggests that an AI-generated score can identify misconduct, or show that an execution met every applicable obligation.

The National Institute of Standards and Technology (NIST) publishes the voluntary AI Risk Management Framework 1.0. [6] It covers responsibilities, tests, monitoring and the retirement of unsuitable uses. We cite the 2023 version. Its resource page lists a revision in progress, and we do not know what the revision will change.

No source in this paper gives a percentage for better trading returns or fewer errors from assistants of this type, and we know of no regional estimate. The evidence supports a bounded test with clear accountability. Test usefulness on the permitted records of the firm, and keep the adverse results.

Evidence review gives no execution authority

The sources establish market infrastructure and conduct duties. None of them measures productivity gains from these proposed assistants.

3 The enterprise lens

Start from a work queue or an approved report that already exists. Define its authoritative sources, and define which timestamp governs each comparison. A daily closing price and an indicative quote answer different questions, and an executable order-book observation answers a third. The assistant must obey one rule: it must not replace one of these with a different one unless the output says so plainly.

Start on read-only snapshots. The assistant gets no credentials for live systems. It cannot place or cancel an order, and it cannot change one. It cannot release funds or amend settlement instructions, and it cannot change limits. Original transaction records stay unchanged. Proposed explanations and corrections go into a separate review record.

Define escalation outside the assistant. If generation fails, existing alerts and settlement deadlines continue, and so do market controls. A material difference that no one can explain goes directly to its designated owner through the established process. Never delay that escalation to wait for a more persuasive explanation. The prompts in the build notes are start points that each firm must test.

Request and source

Assistant output

Reviewer correction

Accepted result

Every explanation must show the original observation with its timestamp and its status. An internal review packet has no path to create an order or change settlement instructions, and no path to suppress an alert.

Check every observation before the desk brief

Business Process Model and Notation (BPMN) swimlanes. Proposed workflow for use case 01.

Desk analyst

AI drafting assistant

Desk research supervisor

Yes

No

Briefing window opens

Select permitted releases and licensed snapshots

Separate facts, scheduled events and questions

Check source times, revisions and restrictions

Brief acceptable?

Accept the internal evidence brief

Desk brief available

Correct stale or unsupported statements

Save the approved internal brief

Pre-market evidence briefing

The assistant produces no trade signals and places no orders. It sends nothing outside the approved audience.

How to read the model. Rounded rectangle: task. X diamond: one decision path. Thin circle: start. Thick circle: end. Solid arrows show the sequence in one process. Lanes show who is responsible.

Notation: Object Management Group, BPMN 2.0.2. https://www.omg.org/spec/BPMN/2.0.2

4 Ten practical use cases

How to read each use case

Each use case has two pages. The first page gives the workflow, the minimum data, the pilot test and the accountable role. The second page gives build notes in five parts.

Starter prompt. Copy the prompt into the approved assistant, and replace the text in square brackets. Each prompt gives the assistant a role, its sources, one task, the rules and the output format.

Tools. The notes name example products. The names are examples and not endorsements. Vendors change product names frequently. Use only a service that the firm approved.

Skills. A staff skill is an ability that the team must build. An agent skill is an instruction file that the assistant loads for the task. Agent Skills is an open format, and many assistants can read it. [8]

Method. The notes name an established method that gives the work a structure.

Agent limits. The notes say what an AI agent can do, and where it must stop.

Four agent levels

Level 0, assist. The assistant answers one request. A person does every step.

Level 1, prepare. The agent collects approved inputs and prepares a draft. It changes no system and sends no message.

Level 2, act with approval. The agent proposes an action. A named person approves each action before it runs.

Level 3, act alone. The agent acts without approval inside fixed limits.

Each pilot in this paper starts at Level 0 or Level 1. A move to a higher level needs a new approval after the pilot. No pilot is at Level 3.

01 Prepare a source-linked pre-market evidence brief

Dated authorised sources

Facts, events and open questions

Desk supervisor checks

Internal pre-market brief

A desk analyst selects the official announcements and the approved economic releases, with licensed market snapshots. The assistant builds a dated brief with four parts: observed facts, scheduled events, open questions and changes since the last approved brief. It recommends no position and predicts no price move.

Each item needs its permitted source, its publication time and observation time, the market calendar and the scope of the desk. A missing timestamp or a release that conflicts goes to review. A delayed price must never look live. An announcement date must never replace the date of the event that it reports.

For the pilot, rebuild a historical morning from only the information that was available then. Put traps into the sources: a revised release and a market holiday, with an old article that someone published again under a new timestamp. Count the unsupported statements and the stale observations, and record the reviewer time. Stop the pilot if restricted information leaks, or if a forecast appears as an official fact.

The desk research supervisor is accountable for the brief. Authorised people judge, under existing controls, if anything in it justifies a trade. The assistant cannot select instruments, issue trade signals, personalise recommendations or send the brief outside its approved audience.

First working session

Rebuild one historical morning brief from information that was available at that hour.

01 Build notes: pre-market evidence

Starter prompt

"You prepare an internal pre-market brief for a desk. Use only the attached sources for [date]. Make four sections: Observed facts, Scheduled events, Open questions and Changes since the last brief. Give the source for each item, with its publication time and observation time. Label a delayed price as 'Delayed'. If a timestamp is missing or two releases conflict, put the item under 'For review'. Do not predict a price move. Do not recommend a position or name an instrument to trade. Do not include information from outside the approved sources."

Tools

Sources: official announcements and approved releases, with a licensed snapshot from the market-data service of the firm, for example Bloomberg.

Brief: Claude or Microsoft 365 Copilot, with the dated sources attached.

Distribution: the approved desk channel, after the supervisor accepts the brief.

Skills

Staff skill: tell the publication time of an item from the time of the event.

Agent skill: a pre-market-brief skill with the four sections and the "Delayed" label.

Method

A point-in-time test: rebuild a past morning from only the information that was available at that hour. Count stale observations and unsupported statements.

Agent limits

Level 1. An agent can assemble the draft brief from the approved sources before the briefing window. It must not send the brief. It has no connection to an order system.

02 Explain market-data quality exceptions

Licensed market snapshot

Approved validation rules

Group data-quality exceptions

Market-data owner investigates

Reviewed exception evidence

Approved checks run first. A market-data analyst applies them to a licensed snapshot. The checks look for missing identifiers, stale timestamps, inconsistent units and duplicate records. The assistant then groups the exceptions and explains which fields need investigation. It does not replace those checks, and it repairs nothing in the source feed.

The minimum data is the instrument reference file, the feed definitions, the usage rights, the event timestamps and the approved validation rules. A missing currency or price convention stops the comparison. A repeated price can be a true price that did not change, or it can be stale. Repetition alone cannot show which.

Test on fictional or licensed historical data that has a unit error and a delayed update, with a value that repeats legitimately. Count the missed defects and the false exceptions, and record the verification time of the analyst. Stop if the assistant invents a replacement price or removes a valid observation. Stop if it uses data that its licence does not cover.

The market-data owner is accountable. Corrections go through the normal controlled process with the approval of the owner. The assistant cannot publish a repaired feed or change valuation inputs, and it cannot clear suspect data for trading. Original values stay on file, each with the evidence behind its proposed explanation.

First working session

Prepare a licensed or fictional snapshot that has one stale timestamp and one price that repeats legitimately.

02 Build notes: market-data exceptions

Starter prompt

"You explain market-data exceptions for an analyst. Approved validation rules already ran on the snapshot. Use only the attached exception list and reference documents. Group the exceptions by rule: Missing identifier, Stale timestamp, Inconsistent unit and Duplicate record. For each exception, show the instrument, the field, the value and the event time. Say what the analyst must check to confirm the cause. Label a repeated price as 'Repeated: cause not known'. Do not suggest a replacement value. Do not remove an observation. Do not use data outside the licence."

Tools

Validation: approved rules in the market-data platform of the firm, or in SQL under change control.

Explanation: Claude or Microsoft 365 Copilot, with the exception list attached.

Exception log: a SharePoint list or a Jira project, with an owner for each exception.

Skills

Staff skill: tell a stale price from a true price that did not change.

Agent skill: a data-exception skill with the four rule groups and the "cause not known" label.

Method

Rule-based validation on four data quality dimensions: completeness, timeliness, validity and uniqueness. Test with seeded errors, and report missed defects and false exceptions separately.

Agent limits

Level 1. An agent can group the exceptions after each scheduled validation run. It must not change or publish feed data. It must not clear data for trading or valuation.

03 Reconstruct trade-exception evidence

Existing trade exception

Order and execution messages

Align observed chronology

Trade-support manager reviews

Exception investigation record

A trade-support analyst selects an exception that the order system or the execution system already raised. The assistant puts the permitted records in time order: orders, acknowledgements, fills and cancellations. It flags missing messages and quantities that disagree. It keeps observed events apart from possible explanations.

Every record needs unique order and execution identifiers, timestamps with a time zone, status definitions and the applicable system logs. Records that the assistant cannot match reliably stay separate. Without supporting records, the chronology must not treat a cancellation request as accepted. It must not treat an acknowledgement as a completed fill.

Use closed, permission-cleared examples together with synthetic failures. Include a partial fill and messages that are out of order. Compare the output with the timeline of an investigator. Count the omitted events and the false links, and count the causes that the records do not support. Stop the test immediately if the assistant proposes a live replacement order or closes an unresolved exception.

The trade-support manager is accountable for the investigation record. Authorised operations staff and traders decide corrections under the procedures of the firm. The assistant cannot cancel or rebook a trade, and it cannot

allocate one. It cannot decide liability, and it cannot treat a plausible technical explanation as an established root cause.

First working session

Select a closed, permission-cleared exception that has a partial fill and messages that are out of order.

03 Build notes: trade-exception chronology

Starter prompt

"You reconstruct the evidence for a trade exception. Use only the attached records for exception [identifier]. Put every message in time order, with its identifier, its type, its status and its timestamp in one time zone. Mark each gap where a message is missing. Mark each pair of quantities that do not agree. Put possible explanations in a separate section with the label 'Not established'. Do not treat a cancellation request as accepted. Do not treat an acknowledgement as a fill. Do not propose a new order or a correction. Do not close the exception."

Tools

Records: order and execution messages from the order management system, for example in the FIX format, with system logs.

Chronology: Claude or Microsoft 365 Copilot, with the exported records attached.

Log search: the log platform of the firm, for example Splunk, used by the analyst.

Skills

Staff skill: read message types and statuses, and put them on one time zone.

Agent skill: a trade-chronology skill with the message types and the "Not established" label.

Method

Event-sequence reconstruction from the message records, ordered by timestamp and sequence number. Compare the result with the timeline of an investigator.

Agent limits

Level 1. An agent can build the chronology for an exception that an analyst selects. It has read access to exported records only. It must not cancel or rebook a trade, and it must not allocate one.

04 Draft a post-trade execution review narrative

Checked execution results

Approved comparison method

Explain costs and limitations

Execution-review lead approves

Post-trade review narrative

In this use case, an execution analyst supplies the inputs: checked fills, costs, client instructions and an approved comparison method. The assistant explains the reviewed calculations and lists their limits, such as partial fills or missing reference observations. The narrative keeps execution quality apart from the later gain or loss of the investment.

The temptation in this use case is to select the benchmark after the desk knows the result. Select it before the review, under the policy of the firm. Assemble the complete execution records, the applicable instructions, the fees and the timestamps. If no one could observe the benchmark at the time, or it does not suit the trade, the draft says so. The draft never implies that a reference price was executable.

Pilot on historical trades whose results a person calculated independently. Include an unfilled remainder and a late data correction. Measure numerical agreement, and count the omitted costs and the unsupported causal claims. Stop if the draft changes the comparison method or hides an adverse result. Stop if it claims that one measure

guarantees best execution.

The execution-review lead is accountable and approves the narrative and its conclusions. The assistant does not recommend the next trade or rank clients for preferential treatment. It does not certify regulatory compliance. Calculations stay in a checked analytical tool with inputs that a reviewer can reproduce.

First working session

Agree the comparison method before anyone looks at the historical execution results.

04 Build notes: execution review narrative

Starter prompt

"You draft a post-trade review narrative from checked calculations. Use only the attached execution results and the comparison method that the firm approved before the review. Do no arithmetic. State the result against the benchmark, with its period and currency, and list the costs. List the limits of the calculation, such as a partial fill or a missing reference observation. Say if no one could trade at the benchmark price. Keep execution quality apart from the later gain or loss of the investment. Do not change the comparison method. Do not hide an adverse result. Do not say that the execution was the best possible. Do not recommend a trade."

Tools

Calculations: the transaction cost analysis (TCA) tool of the firm, or a checked Excel model.

Narrative: Claude or Microsoft 365 Copilot, with the checked results attached.

Method record: the approved benchmark policy, with the date of approval.

Skills

Staff skill: select a benchmark before the result and state its limits.

Agent skill: an execution-review skill with the narrative structure and the forbidden claims.

Method

Transaction cost analysis (TCA) with a benchmark that the firm selects before the review, for example arrival price or volume-weighted average price (VWAP). An independent calculation checks each figure.

Agent limits

Level 1. An agent can draft the narrative when the checked results for a review period are final. It must not calculate or change a result. It must not send the narrative to a client or a regulator.

05 Prepare a desk liquidity scenario discussion

Hypothetical disruption

Market and funding assumptions

Map operational dependencies

Desk and treasury challenge

Reviewed liquidity exercise

A desk supervisor and a treasury specialist define a hypothetical disruption, such as delayed settlement or thinner quoted depth. The assistant organises the approved assumptions and the operational steps that the disruption affects, with the questions for discussion. It shows where a different assumption changes the conversation. It proposes no liquidation and no hedge.

The inputs are the exercise objective, the historical observations, the permitted position aggregates, the settlement dependencies and the assumptions that specialists approved. Missing information on depth or funding stays unknown. A quoted price must never become an estimate of how much the desk could sell at that price.

Pilot with fictional positions and a scenario that combines an execution delay with a funding constraint. Then treasury and operations challenge the dependencies, together with the desk. Count the unsupported assumptions

and the missed operational interactions, and record how far the reviewers disagree. Stop the pilot if someone presents a scenario as a forecast, or if the assistant writes a trading instruction.

The desk liquidity lead is accountable for the exercise. Authorised treasury and trading staff keep the decisions on funding and trades. No simulated result is a price that the desk can achieve. The scope stays with readiness for execution and settlement. Portfolio optimisation and the enterprise stress-testing model belong to other work.

First working session

Write a fictional scenario that contains an execution delay and a funding constraint.

05 Build notes: desk liquidity exercise

Starter prompt

"You prepare a discussion paper for a desk liquidity exercise. Use only the attached exercise objective and approved assumptions. List each assumption with its source and its owner. List the operational steps that the disruption affects, in order. For each step, write one question and name the team that must answer it. Show which assumption changes the discussion most if it is different. Label depth and funding data that are missing as 'Unknown'. Do not estimate a sale price or a quantity. Do not propose a liquidation or a hedge. Do not write a trading instruction."

Tools

Assumptions: a controlled Excel workbook that treasury and the desk approved.

Discussion paper: Claude or Microsoft 365 Copilot, with the workbook attached.

Dependency map: Microsoft Visio or a whiteboard, drawn by the participants.

Skills

Staff skill: state an assumption with its source and challenge it in the room.

Agent skill: a liquidity-exercise skill with the assumption list and the "Unknown" label.

Method

A tabletop exercise with a dependency map and a sensitivity discussion on each disputed assumption. Record the disagreements between reviewers.

Agent limits

Level 0. Specialists run the exercise in a meeting, and a person starts each draft. The assistant has no connection to position systems or funding systems. It must not write a trading instruction.

06 Triage settlement queries without changing instructions

Approved settlement query

Collect authoritative status

Status current and consistent?

Internal query packet

Settlement investigation

Yes

No

A settlement officer supplies an approved query and read-only records. The assistant puts the trade reference, the expected settlement date, the current status and the missing evidence into an internal case note. It uses the definitions of the authoritative system to tell matched instructions from items that are in process, and from confirmed settlement.

The assistant needs verified trade identifiers, the settlement calendar, the current status record and an approved contact directory. A stale source or status terms that conflict send the query to investigation. An email from a

counterparty alone must never override the independently verified instructions of the firm.

Test with fictional queries about a holiday and a failed match, and about an attempt to change bank details. A settlement specialist answers the same set. Compare the classification and the evidence links. Count the incorrect status descriptions and the missed urgent items, and record the preparation time. Stop if the assistant changes payment details or calls an unsettled transaction settled.

The settlement manager is accountable for the queue and for every message to a person. Deadline monitoring operates independently, and required processing never waits for the assistant. The assistant cannot release funds or securities, amend instructions, send confirmations or disclose account details to an unverified recipient.

First working session

Test a fictional query that asks to change the settlement bank details.

06 Build notes: settlement queries

Starter prompt

"You prepare an internal case note for a settlement query. Use only the attached query and the read-only status record. Give the trade reference, the expected settlement date, the current status and the missing evidence. Use the status terms of the settlement system: Matched, In process, Settled or Failed. If the status record is stale or two sources conflict, write 'Send to investigation'. If the query asks to change bank details, write 'Verify independently' at the top and take no other step. Do not call a transaction settled unless the record shows it. Do not draft a confirmation to an external party."

Tools

Status source: the settlement system or the custodian portal, with read access only.

Case note: Claude or Microsoft 365 Copilot, with the query and the status record attached.

Queue: the case tool of the settlement team, for example ServiceNow or Jira.

Skills

Staff skill: verify a change of bank details through an independent, known contact.

Agent skill: a settlement-triage skill with the status terms and the bank-detail rule.

Method

Triage with the status definitions of the authoritative system. Use call-back verification on a known number for any request to change payment details.

Agent limits

Level 1. An agent can prepare a case note for each new query. It has read access only, and deadline monitoring does not depend on it. It must not send a message or change an instruction.

07 Explain existing desk-limit alerts

Existing control alert

Approved limit and exposure

Link versions and caveats

Desk control owner reviews

Alert explanation

The alert comes from the approved limit-monitoring system of the firm, and a control analyst passes it on. The assistant connects the alert to the applicable definition and the checked exposure inputs, with any approved changes that apply. Then it drafts questions about why an apparent movement needs review. All calculation stays in the monitoring system.

The minimum data is the alert identifier, the approved limit version, the calculation timestamp, the exposure basis and the authorised exception record. An absent approval stays absent. A discussion about a possible limit increase gives no approval, even when it is far advanced. Gross exposures and net exposures stay distinct, and so do different currencies.

Use historical or fictional alerts that have a stale limit version and a true breach, with an approved temporary exception. Count the missed caveats and the references to the incorrect version, and record the verification time. Stop if the assistant suppresses an alert or changes its priority. Stop if it calls an unapproved exception authorised.

The desk control manager is accountable for explanation and escalation. Existing controls continue to issue alerts whether the AI operates or not. The assistant sits next to them as an explanation layer. It has no power to raise limits, approve breaches, decrease the monitoring frequency or instruct a trade to resolve an exposure.

First working session

Pair an existing alert with its approved limit version and exposure timestamp.

07 Build notes: limit-alert explanations

Starter prompt

"You explain an existing limit alert for a control analyst. Use only the attached alert record and limit documents. Give the alert identifier and the limit version, with the calculation timestamp. Say if the exposure is gross or net, and give its currency. Check that the limit version was in force at the calculation time. If an exception applies, give its approval record. If no approval record exists, write 'No approved exception'. Write the questions that the control manager must answer. Do no calculation. Do not change the priority of the alert. Do not suggest a trade or a limit increase."

Tools

Alert source: the limit-monitoring system of the firm, with read access.

Explanation: Claude or Microsoft 365 Copilot, with the alert and the limit documents attached.

Escalation: the existing control workflow, which operates without the assistant.

Skills

Staff skill: check that a limit version was in force at the calculation time.

Agent skill: a limit-alert skill with the version check and the exception rule.

Method

A control explanation with a version check and a four-eyes review. The alert and its escalation never depend on the explanation.

Agent limits

Level 1. An agent can draft an explanation for each alert that the system raises. It has read access only. It must not suppress or close an alert, and it must not change its priority.

08 Prepare communication-surveillance review packets

Approved monitoring question

Permitted message context

Context and contrary evidence

Surveillance specialist reviews

Internal review packet

A surveillance specialist selects a sample from authorised business channels. The firm obtained the sample lawfully and approved the access. For human review, the assistant assembles the applicable message excerpts and their context, with linked transaction references. It marks ambiguous language. Evidence against a concern stays

in the packet. The packet gives no verdict on misconduct.

This use case reads staff communications. For that reason, the packet needs an approved purpose, a permitted period, channel records, review criteria and reliable identity mapping. It must flag missing context and speakers that are not certain. Never infer intent or honesty from dialect or style, or from one phrase. Never infer an emotional state.

Pilot on fictional conversations that have quotation, negation, ordinary desk shorthand and some instructions that are a true concern. Write the same meanings in different language styles. Specialists assess the missed concerns and the unnecessary escalations. Stop if the tool exposes unrelated personal material or treats an allegation as a fact. Stop if it silently removes a flagged message.

The surveillance lead is accountable for review and escalation under the compliance process. The assistant cannot clear an investigation, notify an employee, impose discipline or file a regulatory report. It works inside an approved review queue. It gives no authority for covert monitoring or for access to private channels.

First working session

Build fictional conversations that have quotation and negation. Write the same meaning in different language styles.

08 Build notes: surveillance packets

Starter prompt

"You assemble a review packet for a surveillance specialist. Use only the attached message sample and the review criteria. For each flagged message, show the excerpt, the messages before and after it, the channel and the linked transaction reference. Mark language that is ambiguous, and say why. Include evidence that argues against the concern. If the speaker or the context is not certain, say so. Do not infer intent or honesty. Do not infer emotion from style or dialect. Do not include personal material that does not relate to the question. Do not give a verdict. Do not remove a flagged message."

Tools

Sample: an export from the approved surveillance system of the firm, for the approved purpose and period.

Packet: Claude or Microsoft 365 Copilot in a restricted workspace, with access for the surveillance team only.

Review record: the case tool of the compliance function.

Skills

Staff skill: read quotation and negation in context, and know the shorthand of the desk.

Agent skill: a surveillance-packet skill with the review criteria and the list of forbidden inferences.

Method

Two specialists review a sample independently, with fixed criteria. Report missed concerns and unnecessary escalations separately, and test the same meaning in different language styles.

Agent limits

Level 0. A specialist starts each packet from an approved sample. No agent has continuing access to staff communications. The assistant must not notify an employee or file a report.

09 Assemble reconciliation evidence for unresolved breaks

Trade and contract records

Custody and accounting records

Group unresolved differences

Finance and operations verify

Reconciliation evidence

A reconciliation analyst supplies checked comparisons across four record types: trade records, contract notes, custody records and accounting entries. The assistant groups the differences by documented attributes and prepares one evidence packet for each unresolved break. It keeps a timing difference apart from a proposed correction, and it assumes that neither explains the item.

The minimum data is the instrument and transaction identifiers, the quantities, the currencies, the settlement dates, the source-system extracts and the approved matching rules. Any item with missing data stays unmatched. Two errors in separate accounts or instruments can offset exactly. An equal net total must never hide them.

The test set has fictional records with duplicates and partial settlement, and a pair of errors that cancel in aggregate. Check the output against an answer set that a person resolved independently. Count the false matches and the missed breaks, and record the time to investigate their causes. Stop the pilot if the assistant forces a match or invents a balancing entry. Stop the pilot if it closes an open break. We did not test how frequently an assistant finds the offsetting pair.

The reconciliation manager is accountable for clearance. Finance and operations approve their own corrections. The assistant cannot post journals, move assets, overwrite custody data or treat a proposed adjustment as executed. Originals stay on file with the documented reason for each accepted resolution.

First working session

Create two fictional record errors that offset in the total but stay as separate breaks.

09 Build notes: reconciliation evidence

Starter prompt

"You prepare evidence packets for unresolved reconciliation breaks. A fixed-rule comparison already ran. Use only the attached break list and source extracts. Make one packet for each break, at the level of one instrument and one account. Show the values from each record type and name the field that is different. Label a timing difference and a proposed correction as separate items. Do not net two breaks against each other, even if the total is zero. Do not force a match. Do not propose a balancing entry. Do not close a break."

Tools

Comparison: fixed matching rules in the reconciliation system of the firm, or in Excel Power Query.

Packets: Claude or Microsoft 365 Copilot, with the break list and the extracts attached.

Break log: a SharePoint list with the owner and the age of each break, and its approved resolution.

Skills

Staff skill: investigate a break at item level and not at net level.

Agent skill: a break-packet skill with the packet format and the rule "no net between breaks".

Method

Item-level reconciliation with a break log and break ageing. Test with a pair of errors that offset in total, and check that the workflow reports the two breaks.

Agent limits

Level 1. An agent can prepare a packet for each break after the scheduled comparison. It must not post a journal or change custody data. It must not close a break or mark a correction as executed.

10 Replay an operational incident for desk learning

1

Closed event or fictional case

2

Reconstruct known-at-time facts

3

Specialists challenge timeline

4

Team incident discussion

5

Reviewed learning record

An incident owner selects a closed, permission-cleared event, or writes a fictional event. The assistant reconstructs what each team knew at each moment and traces the documented handovers. It drafts questions about missed signals. Later discoveries have their own label. Without that label, hindsight enters the account of decisions that people made at the time.

The inputs are the approved logs, the communications, the incident timeline, the system-clock data and any final findings. Gaps and timestamps that conflict stay visible. A missing record can mean that no action occurred, or that no one logged the action. The logs alone cannot always show which. The absence of a record does not show that someone withheld information deliberately.

Run the replay on a synthetic outage that has delayed messages and an incomplete handover, with a first explanation that was misleading. Compare the replay with an agreed event record. Count the sequence errors and the unsupported attributions, and count the omitted corrective actions. Stop the pilot if the assistant invents a cause. Stop the pilot if it puts blame on an individual without evidence.

The trading-operations incident lead is accountable for the learning review. Specialists put any proposed process change through the existing approval routes. The assistant cannot restart systems or change controls, and it cannot issue disciplinary findings. It looks at past events only. Live incident command and the formal investigation continue as they do without it.

First working session

Select a closed, permission-cleared event. Separate what the team knew at the time from later findings.

10 Build notes: incident replay

Starter prompt

"You reconstruct a closed operational incident for a learning review. Use only the attached incident records. Build a timeline with one line for each event: the time, the team, what the team knew and the source. Put all times on one clock, and show each timestamp conflict. Label facts that the team learned later as 'Known later'. Mark each handover and say what information passed. Where a record is missing, write 'No record' and give no reason for the gap. Then write questions about signals that the team missed. Do not name a cause that the records do not support. Do not blame a person."

Tools

Records: approved logs and communications for the closed event, exported from the log platform of the firm, for example Splunk.

Timeline: Claude or Microsoft 365 Copilot, with the incident records attached.

Learning record: a page in SharePoint, approved by the incident lead.

Skills

Staff skill: separate what a team knew at the time from what it learned later.

Agent skill: an incident-replay skill with the timeline format and the "Known later" label.

Method

A blameless post-incident review, built on a timeline of what each team knew at each moment. Specialists challenge the timeline before the team discussion.

Agent limits

Level 1. An agent can build the draft timeline for a closed event that the incident owner selects. It has no part in live incident command. It must not restart a system or change a control.

Choose the first pilot

These screening scores assume approved read-only snapshots, licensed data, reliable identifiers, qualified desk and operations reviewers, and independent live controls. We did not measure any of the scores. They are our conditional assessments under those assumptions.

Case | Value | Feasible | Data | Effort | Risk

01 Pre-market evidence | 3 | 4 | 4 | Medium | Moderate

02 Market-data exceptions | 4 | 4 | 4 | Medium | Higher

03 Trade-exception chronology | 4 | 3 | 3 | Medium | Higher

04 Execution review narrative | 4 | 3 | 3 | Larger | Higher

05 Desk liquidity exercise | 3 | 3 | 3 | Larger | Higher

06 Settlement queries | 4 | 3 | 3 | Larger | Higher

07 Limit-alert explanations | 3 | 3 | 3 | Larger | Higher

08 Surveillance packets | 3 | 2 | 2 | Larger | Higher

09 Reconciliation evidence | 4 | 3 | 3 | Medium | Higher

10 Incident replay | 4 | 4 | 4 | Medium | Higher

How to read the table

Value is the possible improvement in evidence or the possible decrease in rework. Feasible shows how well the firm can run a controlled test. Data shows how ready the minimum inputs are. Each score is from 1 (weak) to 5 (strong).

Medium effort is a bounded evidence set and one review process. Larger effort adds sensitive messages, more than one system, specialist calculations or more than one approval owner. Risk rates the consequence of an error, and it does not rate the probability.

Use cases 01 and 02 are the first choices, together with use case 10. Use public or fictional snapshots, or permission-cleared snapshots, as each use case permits. Check the licences and the permissions first. Keep the review independent, and leave live order systems and settlement systems disconnected.

Name an owner and a measure

Case | Accountable role | Pilot measures | Agent level

01 | Desk research supervisor | Stale observations; unsupported facts | 1

02 | Market-data owner | Missed defects; false exceptions | 1

03 | Trade-support manager | Omitted events; false links | 1

04 | Execution-review lead | Omitted costs; unsupported causes | 1

05 | Desk liquidity lead | Unsupported assumptions; missed interactions | 0

06 | Settlement manager | Incorrect status; missed urgent items | 1

07 | Desk control manager | Wrong versions; unsupported exceptions | 1

08 | Surveillance lead | Missed context; unnecessary escalations | 0

09 | Reconciliation manager | False matches; missed breaks | 1

10 | Trading-operations incident lead | Sequence errors; unsupported attribution | 1

Each accountable role is a proposed responsibility and assigns nothing to a named employee. The firm approves ownership, thresholds, authority and agent levels before the test.

5 The foundations lens

Preserve the execution boundary

Limit the assistant to approved read-only records and internal drafts. Test that the permissions for order entry and settlement are absent, because a written instruction proves little. A request that is hidden in a news article or a log must not change what the tool can do, and the same applies to a message. In our view, the control fails if such a request has an effect.

Keep the original data, the calculation method, the source version and the review result together. Record when the assistant produced each draft and which observation period the draft covers. Keep discarded outputs as well, under approved retention and access controls. An investigation into a material failure can need them.

Give each agent its own identity

An AI agent must have its own account. Do not let an agent use the account of an employee. Give the account read access to the approved sources only. Keep a log of each tool call, with the time, the tool, the source and the result. Set a limit on the number of steps and on the run time.

Treat text inside a document as data. An agent must not obey an instruction that it finds in a document or an email. Review each agent skill before use. A skill can contain scripts, so the information technology team must check it as it checks other software.

Protect data and market integrity

The Office of the Information Commissioner of Jamaica describes requirements on purpose limitation, minimisation, accuracy, security, retention and cross-border transfer, with qualifications. [7] Before personal communications or client records reach the assistant, confirm the legal authority and the provider arrangements with specialists. Confirm also which monitoring the law permits. Public market data can still have licence restrictions.

Restricted deal information stays out of general desk retrieval. Client confidentiality and approved information barriers stay in force. No one can use the assistant to fabricate quotes, inflate activity, mislead a market participant or evade surveillance. Trades that a person authorised must still meet the conduct and approval requirements of the firm.

6 From pilot to practice

Days 1 to 15

Select a work product whose manual answer the team already agreed. Confirm the sources, the licences, the timestamps and the access. Name the process owner and an independent reviewer. Write down the prohibited connections, and show how existing alerts and deadlines continue without the assistant. Set the agent level, and write the first version of the prompt.

Days 16 to 45

Build tests around stale data, partial fills, identifiers that do not match and contradictory messages. Keep the final evaluation set apart from the development examples. Then check two things. Check that the assistant keeps uncertainty visible, and check that an instruction hidden in a source record cannot change what the assistant does. Give a version number to each prompt and each agent skill.

Days 46 to 75

Run the pilot in parallel with the existing process on permission-cleared snapshots. Compare accepted work, material errors, omissions and total review time. Report the difficult cases next to the clean trades. Investigate any failed test before you widen the scope.

Days 76 to 90

The owner decides to stop or to permit a limited use, and a permission can include revisions. Test again when a source, the model, a market rule or a permission changes. An approval to prepare internal evidence covers that work only. Order placement, communications, settlement and any new investment strategy each need their own authority.

Pilot decision canvas

Complete this canvas with the process owner before you start the assistant. Attach the approved source list and the test results to the decision record.

Desk, market, instruments and work product:

Accountable owner and independent reviewer:

Source systems and their timestamps:

Data licences:

Approved access and information barriers:

Permitted steps and prohibited actions:

Agent level and tool permissions:

Prompt version and agent skill version:

Manual baseline and difficult test cases:

Quality threshold and stop trigger:

Independent alert and deadline controls:

Withdrawal authority and fallback process:

Open dependencies and decision date:

Measure value without invented savings

Compare the total manual time for preparation and review with the assisted time on equivalent work. The assisted time covers preparation, with all verification and correction. Budget for data permissions, integration, tests, specialist review and software. For an agent, add the usage cost of each run. Count the false exceptions and the missed material items, and count each unauthorised disclosure, next to the accepted outputs.

Fewer flagged items in a review packet can mean that the workflow suppressed concerns. Do not read the decrease as better control. Do not credit a short operational pilot with market returns or with trading losses that the desk avoided. The pilot tells you nothing about either. Staff time released is capacity until a real cost change appears. Continue only when the reviewed work meets the agreed standard for quality and safety.

Record | What to enter

Current process | Preparation; review; correction effort; evidence defects

Pilot process | Same measures on comparable cases; all checking included

Costs | Software; data licences; integration; training; supervision; agent usage

Control quality | Accurate status; retained alerts; complete evidence; timely escalation

Decision | Continue (with or without revisions) or stop; owner; review date

References

[1] Caribbean AI Association. Think Tanks, Harbour remit. Accessed 26 October 2026. https://caribbeanaiassociation.com/think-tanks

[2] Bank of Jamaica. Payment System Data Bulletin, May 2025; contents page dated January 2026. Table CYSU05, page 6. https://boj.org.jm/wp-content/uploads/2026/01/PAYMENT-SYSTEM-DATA-BULLETIN-FOR-PUBLICATION-MAY-2025.pdf

[3] Bank of Jamaica. Payment Systems, National Payment System description. Accessed 26 October 2026. https://boj.org.jm/core-functions/financial-system/payment-systems/

[4] Jamaica Stock Exchange. Main Market Rule Book, prevailing-settlement-cycle definition and amendments to rules 311(D) and 312, pages 8, 46–47; effective 27 May 2024. File accessed 26 October 2026. https://cdn.jamstockex.com/pd/2026/07/JSE-Main-Market-Rule-Book-05.pdf

References continued

[5] Financial Services Commission, Jamaica. Market Conduct Guidelines for Securities Dealers, 31 January 2020, sections 4–5 and 8–9. https://www.fscjamaica.org/wp-content/uploads/2020/04/SR-GUID-26-Market-Conduct-Guidelines-Amended-FINAL-FEB-7-2020.pdf

[6] National Institute of Standards and Technology. AI Risk Management Framework 1.0, 2023, section 5, framework core. https://airc.nist.gov/airmf-resources/airmf/5-sec-core/

[7] Office of the Information Commissioner, Jamaica. Data Protection Standards. Accessed 26 October 2026. https://oic.gov.jm/page/data-protection-standards

[8] Agent Skills. Agent Skills Overview. Accessed 10 November 2026. https://agentskills.io

Evidence and version

Evidence reviewed on 26 October 2026. Text revised and build notes added on 11 November 2026. Prepared by the Harbour Think Tank, Caribbean AI Association. Proposed workflows and prompts need local validation before deployment.