Practical AI for Compliance 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 compliance in financial services, one of fourteen business functions in the collection. [1]

We wrote this paper for compliance officers, nominated officers, legal specialists, monitoring teams and the business owners of controls. Nothing in it lets an assistant make a legal determination, clear suspicious activity, impose a sanctions restriction, file a report or certify compliance. Risk management and cybersecurity operations each have their own paper, and so do customer-service follow-ups.

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 compliance checklist must point to the requirement that applies, and it must keep the exceptions with it. I want an investigation summary to show what is still unknown as plainly as what the records establish.

Artificial intelligence can help to compare and retrieve documents, and it can draft text. But an orderly table can look more authoritative than its evidence permits. One dropped cross-reference can misdirect the next decision. So can an outdated list, or a paraphrase that changes an obligation.

I would start small, with a task that a qualified reviewer can check against an authoritative source. Keep the existing controls in operation through the pilot, and make every conclusion traceable to its source.

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. Public sources give the regulatory context and the control concerns. The ten workflows, the screening scores, the prompts and the agent levels are our proposals for local test. They are not findings from a named institution, and no regulator approved them.

An artificial intelligence (AI) assistant retrieves or compares information, and it can draft text. An agent can also do configured steps with connected tools. Every proposal in this paper stops before decisions and external actions. An obligation register records each requirement with its scope, its dates, its owner and its source. A control is any action that the firm takes to meet a requirement or to treat a risk.

Customer due diligence is the work to understand and verify the relevant information about a customer. Anti-money laundering (AML) and countering the financing of terrorism (CFT) usually go together in discussion. A suspicious transaction report records a concern, and it is not proof of a crime. A sanctions match is a possible correspondence that needs assessment. It does not establish identity until that assessment is complete. The legal specialists of the firm confirm which terms and duties apply.

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. [7]

Approved source

Draft or retrieval

Human review

Accepted output

The operating boundary

In every use case, the assistant organises rules and evidence. Authorised officers keep legal interpretation, case decisions, sanctions actions, regulatory filings and compliance conclusions.

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 obligation extraction and regulatory-change comparison, together with approved training scenarios. All of them run on public documents and fictional cases. Customer records and investigation records wait until the firm approves the access and the confidentiality safeguards, and names a qualified reviewer.

Keep the source words next to the interpretation of the assistant, with the entity and the activity that it concerns, and the date. Keep a consultation apart from an operative rule, and keep guidance apart from legislation. Keep an evidence gap apart from a proven breach.

Use the priority table to select a limited pilot. Track four measures: missed conditions, unsupported conclusions, evidence traceability and staff effort after correction. We give little weight to fewer alerts or faster case closure as signs of compliance quality. Nothing in the pilot can delay escalation or a report, or any other action that the law requires.

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 one task and its qualified reviewer. Approve the authoritative sources and the information boundary. Decide what stops the pilot. Decide which existing process takes the work when the assistant is unavailable or incorrect.

Choose a bounded process

Measure the baseline

Test with review

Decide from evidence

1 The regional business case

The Bank of Jamaica published revised guidance in the Jamaica Gazette of 17 June 2026. The guidance covers money laundering and the financing of terrorism, and it covers the financing of proliferation. Its preliminary provisions say which entities it covers and what legal status it has. The resource page of the regulator still has the earlier guidance, which the Bank finalised in 2017 and gazetted in 2018. Label the two versions. Current questions need the new text, and historical questions can need the old text. [2]

A securities dealer must also read the relevant material of the Financial Services Commission (FSC). It cannot assume that Bank of Jamaica guidance settles its obligations. The 2020 Market Conduct Guidelines for Securities Dealers of the FSC cover governance, complaints, communications and continued compliance monitoring. The firm itself must establish which sources govern its actual licence and activities. [3]

Each line of the obligation map of a Caribbean group needs a jurisdiction and a legal entity. A topic heading is too coarse. Familiar words in the rule of another country can still differ in scope or in commencement, or in their exceptions. Keep local legal confirmation visible before anyone proposes a common procedure.

From an authoritative source to checked evidence

We propose this sequence, and we did not measure it anywhere in the region. A qualified specialist confirms applicability. The business owner approves implementation, and the tester examines evidence. An assistant can organise the links, but it cannot change an incomplete chain into a compliance verdict.

Authoritative version

Applicability review

Approved control

Checked evidence

2 What the evidence supports

The Recommendations of the Financial Action Task Force (FATF), in the version updated in June 2026, set international standards. Jamaican legislation still governs. The Recommendations cover customer due diligence, risk assessment for new technologies, suspicious-transaction reports and confidentiality, with other topics. But the obligations of the firm come from local law and regulatory requirements. [4]

In its July 2021 report on new technologies, the FATF sets opportunities next to problems of data quality, explainability, privacy and exclusion. It keeps an important role for human expertise and review. Nothing in that report shows that the generative assistants in this paper decrease false alerts or find more crime. We have no evidence for or against. [5]

Test the hard cases and the incomplete cases. An assistant that retrieves a plausible rule and misses its exception can do more harm than an approved checklist. Compare the assistant with the existing method, and escalate when either method is uncertain.

Trace requirements to checked evidence

This sequence is a proposal, and we did not measure a result. Specialists confirm applicability, and owners approve implementation. Testers examine evidence. The assistant cannot certify the chain.

3 The enterprise lens

Keep four questions apart: what the source says, if it applies to this firm, which control answers it, and what evidence shows that the control operated. Each question needs an accountable reviewer. The output of the assistant must keep them apart where a reviewer can inspect them. No one green status can cover all of them.

Keep a dated source register. For each document, it records the issuer, the document type, the version, the jurisdiction, the relevant entity and the retrieval location. Superseded documents stay for legitimate historical questions, and the register never shows them as current. Log an inaccessible schedule or a missing amendment as an unresolved research issue.

Rule extraction and case determination belong in separate steps. Fixed checks handle arithmetic and required fields, and they do the approved date calculations. Generated text can explain an exception. It must never invent a legal deadline or decide that a transaction is suspicious. Changes to systems and procedures still go through the normal approval arrangements of the firm. The prompts in the build notes are start points that each firm must test.

Request and source

Assistant output

Reviewer correction

Accepted result

Store the source passage, the applicability decision, the control and the operating evidence in separate fields. A reviewer must be able to challenge one field without a silent change to the others.

Retain conditions when interpreting obligations

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

Compliance analyst

AI drafting assistant

Compliance manager and legal reviewer

Yes

No

Authoritative source selected

Provide complete version and entity scope

Extract duties, conditions, dates and exceptions

Determine applicability and resolve interpretation

Interpretation approved?

Obtain owner acceptance of responsibilities

Obligation review complete

Request the missing authoritative source

Interpretation pending

Record approved interpretation and owner

Obligation extraction review

The assistant gives no legal advice to clients and makes no automatic operational rule. It does not approve the register.

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. [7]

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 Extract obligations with their conditions intact

Complete authoritative source

Extract actions and conditions

Compliance and legal review

Draft obligation table

The analyst selects an authoritative document, and the assistant drafts an obligation table from it. Every row has the source passage, the responsible actor, the required action, the condition that starts the duty, the exception and the relevant date. A row also shows any unresolved cross-reference.

The minimum data is the complete source, with its version and status and all its schedules. The firm also supplies a description of its entities. Start with one public guideline. If the assistant cannot find a referenced section, it marks the row incomplete. It fills in nothing from memory or from the text of another jurisdiction.

We expect trouble with conditions, because a draft that drops a qualifier can still look correct. Pilot against a table that a specialist prepared, with many conditional duties and exceptions. Count the missed conditions and the invented obligations, and count the incorrect source locations. Check that no recommendation became an unconditional legal command. Stop if an omitted exception changes a duty. Stop if someone calls an unreviewed interpretation authoritative.

The regulatory compliance manager is accountable, with legal review wherever the text needs interpretation. Specialists decide applicability, and business owners accept implementation. The tool cannot give legal advice to clients or approve the register. It cannot change an extracted sentence into an automatic operational rule.

First working session

Select one complete official guideline. Mark a condition and an exception before the extraction starts.

01 Build notes: obligation extraction

Starter prompt

"You extract obligations from one official document. Use only the attached document, with its schedules. Make one row for each duty. Give the source passage in the exact words, with its section and page. Give the responsible actor, the required action, the condition that starts the duty, the exception and the date. Write 'None stated' where the text gives no condition or exception. If a row refers to a section that you cannot find, write 'Incomplete:

cross-reference missing'. Keep 'must' and 'should' as the source uses them. Do not add a duty from memory or from another jurisdiction. Do not say that a duty applies to the firm. Do not give legal advice."

Tools

Source: the complete official document from the website of the regulator, for example the Bank of Jamaica or the FSC.

Extraction: Claude or Microsoft 365 Copilot, with the complete document attached.

Register: an Excel table or a SharePoint list, with columns for the legal reviewer, the applicability decision, the owner and the review date.

Skills

Staff skill: read a provision and find its condition and its exception.

Agent skill: an obligation-extraction skill with the row format and the cross-reference rule.

Method

Write each obligation as a structured statement with four parts: the actor, the action, the condition and the exception. A specialist compares a sample of rows with the source text.

Agent limits

Level 0. An analyst starts each extraction from a complete source, and legal review follows. The assistant does not write to the obligation register. It must not decide applicability or give advice to a client.

02 Triage regulatory changes by comparing versions

Previous approved text

New official version

Identify changes and dates

Specialists assess impact

Regulatory-change review queue

The assistant compares the approved previous version of an official text with the new version. It lists the changed passages and the affected topics, with the internal documents that possibly need review. A specialist decides how much each change matters.

The assistant needs the two complete versions, each with its issuer and publication date, and its status. It also needs any commencement or transition arrangements that the text states. Check that the two documents cover the same subject and jurisdiction. Publication is not commencement. When the effective date is missing or ambiguous, ask.

Build the pilot set from a known amendment and a consultation, with a change that only alters the format. A useful system finds the substantive change and ignores the format change. Count the missed changes and the cases of false urgency, and record the review time. A consultation that the output describes as binding ends the test. So does a superseded provision that still has the label "current".

The regulatory change lead is accountable. Legal and compliance specialists decide the impact and the timing, and they decide the required action. The assistant prepares the triage queue. Without an authorised decision, it cannot amend policy or notify a regulator, and it cannot set implementation deadlines.

First working session

Get the two official versions. Establish if the new text is operative or a consultation.

02 Build notes: regulatory-change triage

Starter prompt

"You compare two versions of an official text. Use only the two attached versions. First give the issuer and the publication date of each version, and its status, as the documents state them. Then list each changed passage by section, with the old words and the new words. Give each change one label: Substantive, Clarification, Format only or Not sure. Give the commencement date or the transition period only if the text states it. If it does not, write 'Effective date not stated: ask'. List the topics that each substantive change affects. Do not call a consultation a rule. Do not say how the firm must respond. Do not set a deadline."

Tools

Comparison: Word Compare on the two versions, so that every text difference is visible.

Triage: Claude or Microsoft 365 Copilot, with the two versions attached.

Queue: a SharePoint list or Planner, with a specialist as the owner of each change.

Skills

Staff skill: tell publication from commencement, and a consultation from an operative rule.

Agent skill: a change-triage skill with the four change labels and the effective-date rule.

Method

A redline comparison first, then a label for each change. A specialist reviews every change that has the label "Substantive" or "Not sure".

Agent limits

Level 1. An agent can prepare the comparison when the source register receives a new official version. It adds draft entries to the review queue. It must not amend a policy or set a deadline, and it must not contact a regulator.

03 Map approved requirements to policy evidence

Reviewed requirement set

Current policy and procedures

Map provisions and gaps

Policy owner and specialists

Requirement-evidence map

In this use case, the policy owner supplies a reviewed obligation set and the current procedures that answer it. The assistant maps each requirement to a passage or a control description, or to a question about missing evidence. A provision that is absent gets a different flag from one that exists but is possibly too weak.

Include the appendices and the referenced procedures, with the approved requirements, the full policy set, the versions and the control owners. When a requirement depends on an operating instruction that nobody can supply, record the dependency. Leave the requirement open, even if the policy mentions the topic.

Pilot on a bounded policy that has a known omission and a partial response, with a documented exception. A specialist checks every proposed gap. Count the missed material gaps and the false findings, and record the effort to retrieve evidence. The pilot stops if the map treats a matched phrase as proof of compliance. It also stops if a draft drops a stricter approved control without review.

The policy compliance lead is accountable. The business owner proposes changes, and authorised specialists approve them. The output compares documents with documents. It gives no legal opinion and no assurance that staff follow the procedure. Operating effectiveness needs a separate test.

First working session

Select one approved requirement. Find every policy and procedure that you need to assess its coverage.

03 Build notes: policy evidence map

Starter prompt

"You map approved requirements to policy text. Use only the attached requirement set and the attached policies and procedures. For each requirement, quote the policy passage or the control description that answers it, with the document and its version, and the section. Give one result: Present, Partial, Absent or Depends on a missing document. For 'Partial', say which part of the requirement has no answer. Do not treat a matched phrase as an answer. Flag a policy control that is stricter than the requirement, and do not suggest that the firm remove it. Do not say that the firm complies. Do not say that staff follow the procedure."

Tools

Policy set: the policy library in SharePoint, with the current version of each document.

Mapping: a Claude Project that contains only the requirement set and the policy set.

Map: Excel, with one row for each requirement and a column for the specialist decision.

Skills

Staff skill: tell a policy statement from evidence that a control operates.

Agent skill: a policy-map skill with the four results and the matched-phrase rule.

Method

A traceability matrix from each requirement to a policy passage and a control. A specialist checks every "Partial" and "Absent" result before it becomes a gap.

Agent limits

Level 1. An agent can prepare a draft map when a requirement set or a policy version changes. It has read access to the policy library. It must not change a policy or record a gap as confirmed.

04 Practise compliance decisions with fictional scenarios

1

Approved compliance procedure

2

Fictional decision scenario

3

Specialist checks answers

4

Staff practise escalation

5

Unaided understanding check

Take an approved procedure, such as when to escalate missing due-diligence evidence. A training specialist selects it. The assistant drafts a fictional case and a set of possible responses, with an explanation that links to the source. One case must have the correct answer "ask for clarification".

The inputs are the current procedure, a learning objective, the approved escalation route and the marking guide of the specialist. No one copies an identifiable investigation to make a case realistic. Where a rule is missing, the question goes to the trainer. The exercise stays incomplete until someone answers it.

Give staff equivalent tasks before and after practice, and then one task without help. Score the recognition of the relevant conditions and the correct escalations, and record the misunderstandings that recur. Check that protected reporting practices stay inside the authorised training audience. Stop if the exercise rewards tipping-off or false reassurance. Stop if it rewards a person who skips a required control.

The compliance training lead is accountable. Qualified specialists approve scenarios and feedback. Completion and model-generated scores certify the professional competence of no one, and mandatory training still applies. The results improve the teaching. They never go into an automatic disciplinary judgement.

First working session

Write down the approved escalation answer before anyone creates a fictional training case.

04 Build notes: compliance scenarios

Starter prompt

"You draft a fictional practice case for compliance training. Use only the attached procedure and learning objective, with the marking guide. Write one short case with invented names and no real investigation. Write four possible responses. Say which response the procedure supports, and quote the passage. Explain why each other response is incorrect. In one case of every set, make 'Ask for clarification' the correct response. If the procedure does not cover the case, stop and write a question for the trainer. Do not reward a response that warns the customer about a report. Do not reward a response that skips a control. Do not give a score for a person."

Tools

Source: the current approved procedure and the marking guide of the specialist.

Drafts: Claude or Microsoft 365 Copilot, with the procedure attached.

Delivery: Moodle or Microsoft Forms, for the tasks before and after practice.

Skills

Staff skill: recognise the condition that starts an escalation, and use the approved route.

Agent skill: a practice-case skill with the case format and the clarification rule.

Method

Scenario practice with a task before and a task after, then one task without help. This is Kirkpatrick level 2. It measures what staff learned, and it does not measure their later conduct at work.

Agent limits

Level 0. A training specialist starts each case and approves it before use. The assistant has no access to investigation records. It must not record a score against a named person.

05 Prepare balanced transaction-monitoring case summaries

Existing monitoring case

Restricted transaction evidence

Facts and contrary explanations

Investigation lead reviews

Balanced case summary

An investigator supplies an authorised case that the monitoring of the firm already raised. The assistant drafts a chronology and links the relevant transactions. It keeps four things apart: observed facts, explanations received, contradictory evidence and open questions. Information that weakens the initial suspicion stays in the draft.

The minimum data is the restricted case identifier, the transaction extract, the existing alert rationale and the authorised investigation notes. Keep access inside the investigation team. A coherent summary can hide gaps. Missing information on a counterparty or a purpose stays unresolved, and the assistant must not build a criminal narrative to fill the gap.

Run the pilot on synthetic cases and on approved closed files without identities. Investigators compare each summary with the reference record. Include files that the firm closed without suspicion. Count the material omissions and the unsupported allegations, and count the attribution errors. Record the review effort. Stop if the assistant clears an alert or changes its priority. Stop if it shows an unverified explanation as fact.

The nominated officer or an authorised investigation lead is accountable. People decide on investigation and escalation, and on reports, under the requirements that apply. The tool files nothing and changes no detection rule,

and it contacts no customer. A wait for a summary must never delay an existing report or other required action.

First working session

Build a synthetic monitoring case with evidence that supports the initial concern and evidence that weakens it.

05 Build notes: monitoring case summary

Starter prompt

"You draft a summary of one transaction-monitoring case for an investigator. Use only the attached case records. First write a chronology, with the date, the amount, the counterparty and the record for each transaction. Then write four sections: Observed facts, Explanations received, Evidence against the concern and Open questions. Give the alert rationale in the words of the alert. Write 'Not in the records' where information on a counterparty or a purpose is missing. Do not say that activity is suspicious or not suspicious. Do not infer a motive. Do not recommend closure or escalation, and do not recommend a report. Do not draft a message to the customer."

Tools

Case records: the case management system of the firm, with access for the investigation team only.

Summary: Claude or Microsoft 365 Copilot in a restricted workspace that the firm approved for case data.

Test set: synthetic cases, and closed files after de-identification with Presidio or by hand.

Skills

Staff skill: tell an observed fact from an explanation that no one verified.

Agent skill: a case-summary skill with the four sections and the "Not in the records" rule.

Method

A case chronology with four separate sections. A second investigator reviews a sample of summaries against the records, and the sample includes files that the firm closed without suspicion.

Agent limits

Level 0. An investigator starts each summary on one authorised case. No agent has continuing access to case records. The assistant must not clear an alert, change a priority, file a report or contact a customer.

06 Check due-diligence evidence for completeness

Authorised customer file

Current due-diligence checklist

Missing or inconsistent evidence

Quality manager reviews

Verification questions

The assistant gets an authorised customer file and a checklist that compliance approved. It lists evidence that is missing or expired, and evidence that is inconsistent. It links each item to its document. It keeps identity and beneficial ownership as separate questions, and it keeps the purpose of the relationship apart from the two.

The assistant needs the current checklist, the permitted document fields, the evidence dates and the customer category. Use fictional records first. Show an incomplete ownership chain with its missing link visible. Never infer the ultimate owner from a name or an address, or from a web result without support.

Put contradictory dates and an absent ownership link into the test files, with an approved route for alternative documents. The tool must apply the correct checklist, and it must not invent grounds to reject the person. Count the missed gaps and the incorrect expiry flags, and count the unsupported identity assertions. The pilot ends if the tool calls a document authentic or assigns criminal intent. It also ends if the tool treats an incomplete file as permission to open an account.

A due-diligence quality manager is accountable. Verification and acceptance stay with authorised staff under the rules of the firm. The presence of a document verifies nothing. This proposal includes no biometric inference and no new customer risk score. Legal and compliance specialists decide which alternative evidence the firm can accept.

First working session

Put missing ownership evidence into a fictional file, with an approved route for alternative documents.

06 Build notes: due-diligence evidence

Starter prompt

"You check a customer file against a due-diligence checklist. Use only the attached file and the checklist for this customer category. For each checklist item, give one result: Present, Missing, Expired or Inconsistent. Give the document and its date for each result. Report identity and beneficial ownership under separate headings, and report the purpose of the relationship under its own heading. Show each ownership chain link by link, and write 'Link missing' where the file has no evidence. Where the checklist permits an alternative document, say so. Write one verification question for each gap. Do not say that a document is authentic. Do not infer an owner from a name or an address. Do not give a risk score. Do not say that the firm can open or refuse the account."

Tools

File and checklist: the onboarding system of the firm, with the compliance-approved checklist for the customer category.

Check: Claude or Microsoft 365 Copilot in a restricted workspace, on fictional files first.

Question list: the quality-review queue of the due-diligence team.

Skills

Staff skill: follow an ownership chain and find the link that has no evidence.

Agent skill: a file-completeness skill with the four results and the ownership-chain rule.

Method

A checklist review of each file, with a four-eyes check by the quality manager on a sample. Test with seeded gaps, and report missed gaps and incorrect flags separately.

Agent limits

Level 0. A quality reviewer starts each check on an authorised file. The assistant has no connection to public search tools. It must not accept or reject a customer, and it must not contact one.

07 Organise evidence for sanctions-match investigation

Existing potential match

Current approved list entry

Compare identity evidence

Sanctions specialist decides

Investigation evidence pack

The sanctions specialist brings a potential match that the approved screening system already produced. The assistant puts the identifiers and the source-list entries next to each other, with the details that conflict and the evidence that is missing.

The pack always shows the list version and the retrieval time. The inputs are the authorised screening result and the approved list sources, with the permitted identity fields. A similar name does not prove identity. If nobody can confirm that the list is current, the result goes up through the existing sanctions process, and no one clears it.

Of the ten use cases, we are least confident about this one, and its combined screening score is the lowest in the table. Test on fictional records that have similar names and different identifiers, and on incomplete records. Specialists check each draft against reference findings. They look for omitted contradictions and invented identifiers, and for false certainty. Stop if the assistant releases a hold or recommends evasion. Stop if it widens suspicion from a protected trait. Stop if it treats an uncertain match as a confirmed designation.

The sanctions compliance lead is accountable. Only the authorised process decides restrictions, release, escalation and reports. Measures that the law already requires continue without a wait for AI. The assistant cannot freeze or release assets, and it cannot change screening thresholds. It cannot send confidential identifiers to public search tools.

First working session

Use a fictional potential match that has conflicting identifiers and a verified list version.

07 Build notes: sanctions evidence

Starter prompt

"You organise the evidence for one potential sanctions match. The screening system of the firm already produced the match. Use only the attached screening result and list entry, with the customer record. Show the list name and version at the top, with the retrieval time. For each identifier, put the value from the customer record next to the value from the list entry. Give one result for each: Same, Different, Missing in the record or Missing in the list entry. List the details that conflict. If no one confirmed that the list version is current, write 'List not confirmed: escalate'. Do not say that the customer is or is not the listed person. Do not infer identity from a name alone. Do not use nationality or any other protected trait to widen suspicion. Do not recommend a hold or a release. Do not search public sources."

Tools

Match source: the approved sanctions screening system of the firm.

List entry: the approved official list sources, with the version and the retrieval time.

Pack: Claude or Microsoft 365 Copilot in a restricted workspace, with no web access.

Skills

Staff skill: compare identifiers one by one, and confirm that the list is current.

Agent skill: a match-evidence skill with the four identifier results and the list-version rule.

Method

A side-by-side comparison of identifiers, with one of four results for each. A second specialist reviews every pack before a decision.

Agent limits

Level 0. A sanctions specialist starts each pack on one existing match. Legal measures continue without a wait for the assistant. It must not hold or release assets, and it must not change a threshold or clear a match.

08 Classify complaints for compliance review

Approved complaint taxonomy

De-identified complaint text

Multiple issues and uncertainty

Conduct reviewer checks

Compliance oversight themes

Complaints raise compliance themes such as disclosure, charges, conflicts or the handling of instructions. The assistant tags authorised records against the approved conduct-risk classification. It links each tag to the complaint text, and one complaint can have more than one issue.

Start with records without identities. The tool needs the identifier, the source text, the classification definitions and the approved escalation criteria. An allegation stays an allegation. Ambiguous text goes to review. Tone or poor writing is never evidence that a complaint lacks merit.

Pilot with closed cases that contain more than one issue and a rare serious allegation. Compliance staff compare the tags with a reference classification. Look for missed material issues and unsupported tags, and for inconsistent treatment of equivalent words. Stop if the classification downgrades an escalation or suppresses adverse evidence. Stop if it declares a complaint resolved.

The conduct compliance manager is accountable. The complaint owner keeps the handling and the decisions. This use case is for compliance oversight, and service replies and relationship-management follow-ups belong in other papers. For securities dealers, the complaints provisions of the FSC guidelines are a relevant anchor. We would not build a universal deadline engine on them. [3]

First working session

Select complaints without identities that have more than one issue, and include one rare serious allegation.

08 Build notes: complaint classification

Starter prompt

"You tag complaints for compliance oversight. Use only the attached de-identified complaints and the approved classification. Give each complaint every tag that its text supports, and quote the words for each tag. A complaint can have more than one tag. Write 'Ambiguous: review' where the text does not fit a definition. Mark each complaint that meets an approved escalation criterion. Write allegations as allegations. Do not judge if a complaint has merit. Do not use tone or the quality of the writing as evidence. Do not lower an escalation. Do not say that a complaint is resolved."

Tools

Records: an export from the complaint log, after de-identification with Presidio or by hand.

Tags: a Claude Project with the approved classification and its definitions.

Themes: Excel or Power BI, with a count of each tag and a link to the complaints.

Skills

Staff skill: apply the classification definitions, and tag each separate issue in one complaint.

Agent skill: a complaint-tag skill with the classification definitions and the quote rule.

Method

Two reviewers tag a reference set independently, and Cohen's kappa measures their agreement. Then compare the tags of the assistant with the reference set.

Agent limits

Level 0. A conduct reviewer starts each batch of de-identified complaints. The complaint owner keeps every handling decision. The assistant must not close a complaint or lower an escalation.

09 Check a regulatory reporting pack before submission

Draft regulatory return

Instructions and source totals

Arithmetic and field checks

Designated signatory reviews

Pre-submission evidence pack

An authorised reporting officer supplies a draft return and its approved instructions. Fixed checks test the required fields, the reporting periods, the units and the arithmetic. They also test agreement with signed source totals. The

assistant explains each exception and links it to the field and the instruction.

The assistant needs the current return specification, the reporting entity, the period, the validated source data and the approved reconciliation rules. If an instruction or a source total is missing, the item stays unresolved. No one fills an empty field with a statistically plausible number, and no one assumes that it means zero.

Give the checks a fictional return that has a unit error and an inconsistent period, with a missing schedule. The reporting team does the checks again independently. Count the seeded errors that the checks missed and the false exceptions, and record the correction time. Stop if the assistant changes a certified total. Stop if it removes an exception to make the return balance, or submits anything externally.

The regulatory reporting compliance lead is accountable. The designated signatory approves the complete pack, and authorised staff submit it. A pass on format checks does not establish that a return is accurate and complete, or that it is legally sufficient. We did not measure how frequently a return passes those checks and fails the substantive requirements.

First working session

Create a fictional return that has a unit error and a missing schedule, with independent source totals.

09 Build notes: reporting quality checks

Starter prompt

"You explain exceptions in a draft regulatory return. Fixed checks already ran. Use only the attached exception list and return instructions, with the signed source totals. Group the exceptions by check: Required field, Period, Unit, Arithmetic and Source total. For each exception, give the field and its instruction, with the two values that do not agree. Say what the reporting officer must check. Write 'Unresolved' where an instruction or a source total is missing. Do no arithmetic. Do not suggest a value for an empty field. Do not read an empty field as zero. Do not change a total. Do not say that the return is accurate or complete."

Tools

Checks: fixed rules in Excel Power Query, or in the reporting tool of the firm, under change control.

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

Evidence pack: a SharePoint folder with the checks and the exceptions, and the sign-off of the signatory.

Skills

Staff skill: tie a figure in the return to its signed source total.

Agent skill: a return-exception skill with the five check groups and the empty-field rule.

Method

Fixed validation rules, then a tie-out of each total to its signed source. Test with seeded errors, and report missed errors and false exceptions separately.

Agent limits

Level 1. An agent can explain the exceptions after each run of the fixed checks. It has read access to the draft return. It must not change a figure or submit a return.

10 Prepare evidence for a compliance test

Approved compliance test plan

Selected sample evidence

Separate rule and observation

Independent tester determines

Draft test working papers

The tester supplies an approved objective and a population definition, with the sample method. The assistant then organises the selected evidence and drafts working-paper questions. It keeps the rule, the test step, the observation and the proposed finding in separate fields.

The inputs are the approved test plan, the authorised sample, the relevant procedure and the evidence dates. The tester draws the sample by the agreed method. Records that no one can find stay as exceptions. The assistant never replaces difficult items with easier ones to complete the file.

Use a completed test that has known passes and failures, with unresolved items. A second tester reviews the draft without sight of its proposed conclusion. Count the unsupported findings and the missed exceptions. Measure source traceability and review effort also. Stop if the assistant changes the sample or invents evidence. Stop if it describes a limited test as assurance over the whole business.

Qualified testers decide findings and challenge management responses, and they approve closure. The compliance monitoring manager is accountable. Independence stays intact. The assistant never validates its own work, and it never certifies remediation because a new document arrived.

First working session

Approve the test objective and the sample before you ask the assistant to organise evidence.

10 Build notes: compliance test evidence

Starter prompt

"You organise the evidence for a compliance test. Use only the attached test plan and procedure, with the evidence for the sample that the tester selected. Make one working-paper row for each sample item, with four separate fields: Rule, Test step, Observation and Proposed finding. In 'Observation', state only what the evidence shows, with the document and the date. Write 'Evidence not found' where a record is missing, and keep the item in the sample. Label each proposed finding 'Draft: tester decides'. Do not replace or remove a sample item. Do not invent evidence. Do not describe the test as assurance over the business."

Tools

Plan and sample: the approved test plan, and the sample that the tester draws in Excel.

Working papers: Claude or Microsoft 365 Copilot for the draft rows, and Word or Excel for the tester.

Evidence: a SharePoint folder for the test, with one reference for each sample item.

Skills

Staff skill: keep an observation apart from a finding, and keep every sample item.

Agent skill: a working-paper skill with the four fields and the "Evidence not found" rule.

Method

An attribute test on a sample, with working papers in four fields. A second tester does a sample of the steps again without sight of the proposed finding.

Agent limits

Level 0. The tester starts each draft after the sample is final. The assistant has no part in sample selection. It must not record a finding or close an item, and it must not certify remediation.

Choose the first pilot

These screening scores assume a firm that has current official sources and qualified compliance reviewers, with restricted workspaces. They also assume no autonomous determinations and no external actions. The scores are our conditional first assessments, and no one measured them.

Case | Value | Feasible | Data | Effort | Risk

01 Obligation extraction | 4 | 4 | 4 | Medium | Higher

02 Regulatory-change triage | 4 | 4 | 4 | Medium | Higher

03 Policy evidence map | 4 | 3 | 3 | Medium | Higher

04 Compliance scenarios | 3 | 4 | 4 | Medium | Moderate

05 Monitoring case summary | 4 | 3 | 2 | Larger | Higher

06 Due-diligence evidence | 4 | 3 | 3 | Medium | Higher

07 Sanctions evidence | 3 | 2 | 2 | Larger | Higher

08 Complaint classification | 3 | 4 | 3 | Medium | Higher

09 Reporting quality checks | 4 | 3 | 3 | Medium | Higher

10 Compliance test evidence | 4 | 3 | 3 | Medium | Higher

How to read the table

Value is the possible gain in evidence quality or the 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 includes specialist interpretation or evidence checks, or assessed training. Larger effort needs protected investigation records and current screening evidence, with tighter access tests. Risk rates the consequence of an error only.

Use cases 01 and 02 are the first choices, together with use case 04. Use complete public documents and fictional examples. Keep investigation records and due-diligence records out of the first pilot, and keep sanctions records out also.

Name an owner and a measure

Case | Accountable role | Pilot measures | Agent level

01 | Regulatory compliance manager | Missed conditions; invented obligations | 0

02 | Regulatory change lead | Missed changes; false urgency | 1

03 | Policy compliance lead | Missed gaps; false findings | 1

04 | Compliance training lead | Correct escalation; recurring errors | 0

05 | Nominated officer | Material omissions; unsupported allegations | 0

06 | Due-diligence quality manager | Missed gaps; unsupported identity claims | 0

07 | Sanctions compliance lead | Omitted contradictions; invented identifiers | 0

08 | Conduct compliance manager | Missed issues; inconsistent classification | 0

09 | Regulatory reporting compliance lead | Missed errors; false exceptions | 1

10 | Compliance monitoring manager | Unsupported findings; missed exceptions | 0

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

Keep judgement and protected reporting with authorised officers

An assistant has no legal authority, even when its text is confident. Keep applicability, breach findings, customer acceptance, sanctions decisions and regulatory filings with the designated people. Reviewers need the original sources and the relevant competence. They also need the authority to reject a draft. A duty to click "approve" is not enough.

Keep investigative material out of general business knowledge stores. Protect suspicious-report information and the related deliberations from unauthorised disclosure, and keep them out of customer communications. The FATF

standards cover the confidentiality of reports. The specialists of the firm must confirm the local duties and the permitted arrangements to share information. [4]

Access to service must never depend on fluency or on inferred personality. It must not depend on the impression that an assistant has of credibility. Test equivalent cases with different words and with incomplete evidence. Correct errors through a documented human route that keeps lawful restrictions and required controls in force.

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.

Maintain lawful access and a dependable fallback

The Office of the Information Commissioner of Jamaica gives standards on purpose, minimisation, accuracy, retention, security and transfers outside Jamaica. Compliance records stay personal data wherever they can identify a person. For each use and each external provider, the responsible specialists must confirm the lawful arrangements. [6]

Use restricted workspaces and read-only connections. Before any record goes in, approve the storage of the provider, its subcontractors, its retention terms and its model-training terms. Instructions hidden in retrieved files are untrusted content. They have no authority to disclose information or to call more tools.

Keep source versions and material corrections under the authorised retention policy, with the approval history. Test isolation between cases and test deletion. Check that records under legal hold survive. Test again after a model update or a changed source library, and after a new integration. The established process must stay available through outages and refusals, and through investigations of faulty output.

6 From pilot to practice

Days 1 to 15

Select a task on public documents, and name its qualified reviewer. Measure the current handling time and the types of error that occur now. Then approve the source versions, the permitted actions, the reference answers and the person who can stop the pilot. Put exceptions and missing cross-references into the test set. Set the agent level, and write the first version of the prompt.

Days 16 to 45

Run the assistant next to the existing workflow, and check every output. Log the origin of each error: source, retrieval, interpretation or calculation. Keep escalation and reporting deadlines outside the control of the pilot. Give a version number to each prompt and each agent skill.

Days 46 to 75

Compare accepted outputs on equivalent work, with review and correction counted. Test failure conditions with obsolete sources and inaccessible schedules. Simulate a provider outage. Write equivalent cases in different words, and check that the treatment stays consistent. Correct the process before more sensitive information comes in.

Days 76 to 90

The sponsor reviews material errors, confidentiality tests, unresolved interpretations, staff effort and cost. The sponsor decides to continue or to stop, and a decision to continue can include revisions. Record the monitoring owner and the fallback route. A new jurisdiction or dataset needs its own approval, and so does a new action.

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.

Process and accountable officer:

Requirement or evidence problem:

Jurisdiction and entity:

Source version:

Permitted data and authorised reviewers:

Actions the assistant may perform:

Actions it must never perform:

Agent level and tool permissions:

Prompt version and agent skill version:

Baseline and independently checked test cases:

Material-error threshold and review method:

Stop condition and existing fallback process:

Decision date and evidence required:

Measure value without inventing savings

Count work that reviewers accepted as correct, and include the time for research, review, correction and administration. A harmless change of words goes in one column. An omission that could change the decision of an officer goes in another column. Difficult cases and unresolved questions stay in the comparison, even when they make the pilot look worse.

Include the costs of source maintenance, licences, integration, training and supervision. For an agent, add the usage cost of each run. Released time is capacity, and it counts as a cash saving only when the firm decreases expenditure. Fewer reports or alerts can mean weaker scrutiny, and so can faster closures. Read them against case quality that someone reviewed independently.

A short pilot cannot show that financial crime fell or that enforcement risk went away, and we claim neither. The decision record must show the sample size, the consequences of errors, the corrected examples and the review burden. If the system breaks confidentiality or distorts a material obligation, stop the pilot, even when the average handling time improved.

Record | What to enter

Current process | Work volume; handling time; review time; correction rate

Pilot process | Same measures on comparable work; all review and correction effort

Costs | Licence; integration; training; continuing supervision; agent usage

Compliance quality | Missed conditions; unsupported findings; confidentiality failures

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

References

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

[2] Bank of Jamaica. 2026. Guidance Notes on the Prevention of Money Laundering and Countering the Financing of Terrorism, Proliferation and Managing Related Risks. Jamaica Gazette Extraordinary, 17 June, No. 277; preliminary provisions, pages 915–917. Current and earlier editions listed at the official anti-money-laundering resource page, accessed 26 October 2026.

https://boj.org.jm/core-functions/financial-system/boj-anti-money-laundering-aml/

[3] Financial Services Commission, Jamaica. 2020. Market Conduct Guidelines for Securities Dealers, SR-GUID-20/01-0026, 31 January. Sections 3, 7, 8 and 14; effective 1 May 2020. https://www.fscjamaica.org/wp-content/uploads/2020/04/SR-GUID-26-Market-Conduct-Guidelines-Amended-FINAL-FEB-7-2020.pdf

[4] Financial Action Task Force. International Standards on Combating Money Laundering and the Financing of Terrorism & Proliferation. Updated June 2026; Recommendations 6–7, 10, 15 and 20–21. https://www.fatf-gafi.org/content/dam/fatf-gafi/recommendations/fatf-recommendations-2012.pdf

References continued

[5] Financial Action Task Force. 2021. Opportunities and Challenges of New Technologies for AML/CFT (anti-money laundering and countering terrorist financing). July; executive summary, sections 4.2–4.3. https://www.fatf-gafi.org/content/dam/fatf-gafi/guidance/Opportunities-Challenges-of-New-Technologies-for-AML-CFT.pdf

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

[7] 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.