CARIBBEAN AI ASSOCIATION / HARBOUR THINK TANK

Practical AI for Sales 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 is one of fourteen in the collection. It looks at the work around a sale in financial services: to understand an enquiry, to find the approved information, to prepare an offer and to keep an accurate record of what occurred. [1]

We wrote this paper for sales managers, product teams, marketing staff, compliance specialists and business sponsors. Investment suitability and individual advice belong to qualified people. The relationship-management paper covers continuing service, and the market-research paper covers market sizing and customer surveys.

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 client who asks about an offer needs the conditions together with the headline benefit. The salesperson needs to know which version of the product information is current, and what the firm has authority to promise. I count more persuasive wording as a loss when it hides a charge or creates an expectation that the firm cannot meet.

An assistant can help to prepare the comparison and collect the documents, and it can show what is missing. It must also make a decision of the client not to proceed easier to accept. For that reason, never judge sales quality from the conversion rate alone.

Start with an internal task whose answer the team can check. Keep product approval, client eligibility, advice and permission to contact as four separate decisions. A tool that drafts well has no authority to make any of them.

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 institutions that we cite supply context and control requirements, and none of them endorses this paper. The ten workflows, the priority scores, the prompts and the agent levels are our proposals for tests. They are not results from the sales operation of a named firm.

In this paper, artificial intelligence (AI) is software that predicts or classifies, and that can also generate content. An assistant prepares work for review. An agent can also do configured steps with tools. No proposal in this paper authorises autonomous outreach, recommendations, pricing changes or the acceptance of an application from a client.

An enquiry is a request for information. It does not prove that a person wants to buy or is eligible to buy. A sales pipeline records opportunities and the evidence behind each stage. A conversion rate needs a defined start group and a defined result, with a measurement period. A suppression list contains the contacts that the firm excludes from one communication purpose. Whether an approved product suits a given client is a separate question, and the product document does not answer it.

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. [5]

Approved source

Draft or retrieval

Human review

Accepted output

The operating boundary

The assistant prepares product facts and draft materials, and it builds review queues. Authorised people keep authority over communications, eligibility, advice, pricing and official records.

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 approved product comparisons and proposal assembly, then add pipeline checks. The firm can test each of them internally. No test sends a message or ranks a person by how likely that person is to buy. Contact lists come only after the firm verifies the purpose and the permission, and tests the suppression controls.

Keep four questions apart: what the person asked for, what the firm can offer, what evidence supports the terms and who can approve the next step. Missing information must lead to a question. The assistant must never fill the gap with a guessed income or eligibility result, or a guessed investment preference.

Measure incorrect claims, omitted conditions, review effort and accepted handovers together with any commercial result. Higher conversion does not excuse a misleading explanation or unwanted contact. Before you credit the assistant with a result, compare equivalent opportunities. Allow for changes in products and channels, and in the customer mix.

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.

Decisions for the sponsor

Select one product family and one internal workflow. Name the product owner and the reviewer. Approve the source versions and the prohibited actions, and agree which failure stops the pilot. Keep revenue targets out of quality-assessment decisions, so that the targets do not encourage reviewers to overlook errors.

Choose a bounded process

Measure the baseline

Test with review

Decide from evidence

1 The regional business case

A Caribbean firm can sell through more than one legal entity, in different currencies and under different permissions. Two products with the same label can then have different terms or availability. Before you prepare a comparison, record the entity, the jurisdiction, the currency, the version and the intended audience.

The Jamaican guidelines for securities dealers ask for communication that clients can understand, and for disclosure of terms before an agreement. [2] Deposit-taking institutions have their own customer-related code, which covers terms and fees, and the disclosure of interest rates. [3] The scopes are different. Neither document is a rulebook for every financial business, and the compliance owner must confirm which obligations apply.

Sales records contain personal information too. The data protection standards of Jamaica cover permitted purpose, minimisation, accuracy and security. [4] The fact that a firm has the email address of a person gives no permission to put that address on a new marketing list. Send the proposed use and the channel arrangements to the responsible specialists for confirmation first.

The diagram shows a proposed local evidence chain. We have no regional benchmark for conversion, so the firm must define its own start group and stages before it reports a percentage. Public financial-inclusion data does not show two things that a sales team needs to know. It does not show if the firm had authority for later contact, and it does not show if the offers were accurate.

Record the request

Check contact purpose

Verify offer terms

Record the decision

2 What the evidence supports

The 2020 Market Conduct Guidelines for Securities Dealers of the Financial Services Commission (FSC) cover clear communication that does not mislead. [2] They also cover disclosure of terms and current product information. Section 8.2(iv) covers a no-objection letter for new marketing materials for new products and services, before public dissemination. An internal AI check cannot replace the approval process that applies.

Article 8 of the 2016 customer-related code for deposit-taking institutions in Jamaica requires clear terms and applicable disclosures. [3] The disclosures cover fees and the presentation of interest rates, with effective annual rates among them. A comparison that mixes rates on different bases, or omits a condition, can need more than an edit of the words. Do not rely on a language model to make rates comparable. Use checked source values and validated calculations, and ask a specialist to confirm the basis.

These sources support disciplined preparation and disclosure. They do not measure the effect of an AI tool on sales. On this evidence, no one knows if personalised persuasion improves results for clients. The use cases test the quality of the work first. A claimed commercial gain needs its own comparison design, and you must read it together with complaints, cancellations and reports of unwanted contact.

Check the offer as well as the conversion

A conversion rate can increase while disclosure is inaccurate or a recommendation is unsuitable. Review the errors and the unwanted contact, and look at the result for the client.

3 The enterprise lens

Map an enquiry from its arrival to an accepted handover or a recorded decision not to proceed. Record the origin of the request, the current product information, the permission to communicate and the human approval. A polished proposal must not hide a missing step.

Start with a dated source library and fictional enquiries. Keep the draft of the assistant apart from the official opportunity record. The assistant must not scrape private profiles or infer vulnerability. It must not create urgency from the financial circumstances of a person.

The ten use cases are bounded preparation tasks, and the prompts in the build notes are start points that each firm must test. If an agent later sends or updates anything, that action needs its own approved workflow. The workflow needs recipient checks, duplicate prevention, logs and a way to reconcile uncertain results. A successful internal pilot gives none of those permissions.

Request and source

Assistant output

Reviewer correction

Accepted result

Keep the request and the approved offer traceable, with the human decision. Missing permission or evidence must hold the next step. No one can fill the gap with a guessed profile or an automatic message.

Route enquiries without inferring eligibility

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

Sales employee

AI drafting assistant

Sales operations lead

Yes

No

Enquiry received

Select enquiry and approved routing rules

Propose queue; flag unclear request

Check service scope and appropriate route

Route supported?

Accept the internal referral route

Routing complete

Clarify request or refer

Awaiting resolution

Record the accepted referral route

Enquiry routing

The assistant must not infer wealth or decide eligibility. It must not assess suitability or open a new contact channel.

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. [5]

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 Route an enquiry without inventing a client profile

Authorised enquiry

Propose appropriate queue

Request clear and in scope?

Employee-confirmed referral

Clarify or refer to specialist

Yes

No

The assistant summarises what the person asked for. It decides if the request is for information or is a question about an application. A request for individual advice is a third type. Then it proposes a queue, which is product information or a specialist. A sales employee confirms the route.

If the product or the jurisdiction is unclear, the assistant drafts a clarification question. It must not guess from a name or an address, and the pilot must catch any guess of that type. The assistant needs only the authorised enquiry and the current service boundaries, with the routing rules. A request for a general brochure never justifies a full financial profile.

Pilot on fictional enquiries. Make some of them ambiguous, and include some that the firm cannot serve. Count the incorrect routes and the unnecessary data requests, and record the time to an accepted referral. Stop the pilot if the assistant denies access to a service or invents eligibility. Stop the pilot if it ranks a person by inferred wealth or sensitive traits.

The sales operations lead is accountable. A route from this use case decides nothing about suitability or credit, and it gives no one authority to contact the person through a new channel.

First working session

Select one fictional enquiry. Find the information that you need to route it without a profile of the person.

01 Build notes: enquiry routing

Starter prompt

"You route an enquiry for a sales team. Use only the attached enquiry and the routing rules. Summarise the request in two sentences. Name the request type from the routing rules. Propose one queue and give the rule that supports it. If the product or the jurisdiction is not clear, write one clarification question and propose no queue. Do not

guess from a name or an address. Do not decide eligibility or suitability. Do not rank the person. Do not ask for data that the route does not need."

Tools

Enquiry log: the CRM of the firm, for example Salesforce or Microsoft Dynamics 365.

Routing draft: Microsoft 365 Copilot, or Claude with the enquiry and the rules attached.

Routing rules: a controlled document in SharePoint, with a version number.

Skills

Staff skill: write a clarification question that asks for the minimum data.

Agent skill: an enquiry-routing skill with the request types and the list of forbidden inferences.

Method

Rule-based triage with a human check. Test with a set of fictional enquiries and a confusion matrix: correct routes against incorrect routes for each queue.

Agent limits

Level 1. An agent can propose a queue for each new enquiry. It must not move the enquiry or reply to the person. A sales employee confirms each route.

02 Compare approved product facts on the same basis

Approved product A facts

Approved product B facts

Compare on common basis

Product specialist checks

Factual comparison

The output of this use case is a factual comparison from current approved product documents. Fees, currency, term, access conditions, material risks and limitations each get a row and a source. The comparison does not select a product for the reader.

The inputs are complete product documents with effective dates, and an agreed comparison structure. If units or rate conventions are different, a person must check the calculations. The comparison marks a missing condition as a gap, and it does the same for figures that are not fairly comparable. It puts no rank on a gap.

For the pilot, select products that have conditional fees and different currencies, with different access restrictions. A product specialist checks every field, and a second reviewer checks that the layout does not hide a material difference. Count the unsupported entries and the omitted conditions. Stop the pilot if the comparison invents a guarantee or becomes a personalised recommendation.

The product owner is accountable. The comparison supports a conversation. Individual suitability belongs to the authorised adviser, and any external use needs the approval of the firm.

First working session

Compare two approved product documents. Mark each condition that prevents a like-for-like comparison.

02 Build notes: product comparison

Starter prompt

"You compare two approved products on the same basis. Use only the attached product documents [titles and effective dates]. Make a table with one row for each of these: fees, currency, term, access conditions, material risks and limitations. Give the page or the clause for each cell. If a rate uses a different convention, write 'Not comparable: specialist to calculate'. If a condition is missing, write 'Gap'. Do no arithmetic on rates. Do not rank the products. Do not say which product fits a person."

Tools

Product documents: the approved product library in SharePoint, with effective dates.

Comparison: Microsoft 365 Copilot in Word or Excel, or Claude with the two documents attached.

Rate calculations: a validated Excel workbook that a specialist owns.

Skills

Staff skill: put two rates on the same basis and state the convention.

Agent skill: a product-comparison skill with the row template and the gap rule.

Method

Like-for-like comparison with one source for each cell. A second reviewer does a prominence check: no material difference sits where a reader can miss it.

Agent limits

Level 1. An agent can prepare a draft comparison when a product document changes. It must not publish the comparison or send it to a client. The product owner approves each version.

03 Assemble a proposal from approved components

Approved terms and templates

Assemble versioned pack

Proposal manager checks

Draft proposal pack

We treat proposal assembly as a version-control task. The assistant builds a draft pack from approved descriptions, fee information, disclosures and application instructions. It records which version supplied each section and which approvals are missing. The pack stays a draft until a person checks it.

The inputs are the permitted product set, the approved templates, the current terms and the stated request of the person. Use fictional client details at the start. If an attachment or an approval is missing, the pack stays marked incomplete. No one can take terms from a different offer to close the gap.

Test the workflow on completed example packs with known required documents. Include an expired offer and a conditional charge. Compare the omissions and the inconsistent figures, and record the preparation time with review. Stop if the draft changes a contractual term or invents a discount. Stop if it implies that the firm accepted an application.

The sales proposal manager is accountable. The product and legal approvals keep their existing processes, and so does the compliance approval. A complete pack says nothing about whether the product suits its recipient.

First working session

Assemble a fictional proposal with one attachment missing. Check that it stays marked incomplete.

03 Build notes: proposal assembly

Starter prompt

"You assemble a draft proposal pack. Use only the approved components in the attached list. For each section, give the component name and its version. List each required attachment and approval, and mark it Present or Missing. If one item is missing, put 'INCOMPLETE DRAFT' at the top of the pack. Do not write new terms. Do not change a fee or a date. Do not take text from a different offer. Do not say that the firm accepted an application."

Tools

Components: an approved content library in SharePoint, with a version for each component.

Assembly: Microsoft 365 Copilot in Word, or Claude with the component list attached.

Approval record: the existing approval workflow of the firm, for example in Power Automate.

Skills

Staff skill: check a pack against the required-document list.

Agent skill: a proposal-assembly skill with the section order and the "incomplete" rule.

Method

Document assembly from a controlled component library, with a checklist of required documents for each product. Compare each draft with a completed reference pack.

Agent limits

Level 1. An agent can assemble a draft pack when a salesperson asks for one. It must not send the pack. It must not record an approval that a person did not give.

04 Check the pipeline against evidence

Opportunity records

Approved stage definitions

Flag unsupported stages

Record owner reviews

Proposed corrections

The assistant checks opportunity records against the approved stage definitions. It flags a stage that has no evidence, and an expected date that is in the past. It also flags two entries that can be the same opportunity. Corrections go to the record owner as proposals.

The minimum data is authorised opportunity metadata, stage rules, dates and source references. Where a dated status record is sufficient, do not copy the full client conversation. Record a missing response as unknown. It is not an agreement, and it is not a rejection.

For the pilot, use an extract that a person checked by hand. Add stage errors and entries that look like duplicates. Count the unsupported stages and the incorrect flags, and record the correction effort. Stop if the assistant changes a record or merges opportunities. Stop if it invents a probability that a sale will close.

Every update needs the approval of the owner, and the sales operations manager is accountable. Keep each original entry and the reason for its correction, so that a cleaner pipeline still shows the history behind a target.

First working session

Check five opportunities against their stage evidence. Keep the original records unchanged.

04 Build notes: pipeline evidence

Starter prompt

"You check a sales pipeline extract against the attached stage definitions. For each opportunity, say if the record has the evidence that its stage needs. Flag each expected date that is in the past. Flag each pair of entries that look like the same opportunity, and give the reason. Record a missing response as 'Unknown'. Write each correction as a proposal for the record owner. Do not change a record. Do not merge entries. Do not estimate a probability of sale."

Tools

Extract: the pipeline report from the CRM, with metadata only.

Check: Excel or Power BI for dates and duplicates, and Claude or Microsoft 365 Copilot for the evidence test.

Proposals: a review list in SharePoint or the CRM for the record owner.

Skills

Staff skill: state the evidence that each stage needs.

Agent skill: a pipeline-check skill with the stage definitions and the rule "propose, do not change".

Method

Seeded-error test on a hand-checked extract. Report the missed errors and the false flags separately, and keep an audit trail of each correction.

Agent limits

Level 1. An agent can check the pipeline extract each week and prepare proposals. It must not update a stage or a date. It must not merge or close an opportunity.

05 Prepare for a requested product-information call

Agreed information request

Facts and specialist questions

Sales team lead checks

Requested-call brief

The assistant prepares a short brief for a conversation that the person agreed to. The brief has the enquiry, the topics that the person raised, the current product facts and the questions that only a specialist can answer. Each fact has a link to an approved source. The brief contains no personality profile of the person.

The minimum data is the recorded request, the agreed contact purpose, the permitted notes and the current product material. Do not import unrelated browsing history or observations from social media. If the contact arrangement or the product version is unclear, settle it before the call.

Pilot with fictional enquiries and a reference brief that a reviewer wrote. Include one question that the salesperson is not permitted to answer. Count the unsupported claims and the missed questions, and record the preparation time. Stop if the brief suggests that the salesperson exploit financial distress or hide a fee. Stop if it treats an enquiry as permission for unrelated promotion.

The sales team lead is accountable. This use case prepares a call that a person asked for. It does not start calls or recommend investments, and the full review by the adviser stays in place.

First working session

Prepare a brief for one agreed fictional call. Include a question that needs an adviser.

05 Build notes: call preparation

Starter prompt

"You prepare a brief for a call that the person asked for. Use only the attached request record and product material. List the topics that the person raised. Give the current product fact for each topic, with its source and effective date. List the questions that only a specialist or an adviser can answer. Do not describe the personality or the finances of the person. Do not suggest a way to create urgency. Do not add a product that the person did not ask about."

Tools

Request record: the CRM of the firm, with the agreed contact purpose.

Brief: Microsoft 365 Copilot in Word, or Claude with the request and the product material attached.

Product facts: the approved product library, current version only.

Skills

Staff skill: tell a question for sales from a question for an adviser.

Agent skill: a call-brief skill with the brief sections and the forbidden content.

Method

Compare each brief with a reference brief from a reviewer. Count unsupported claims and missed questions separately.

Agent limits

Level 1. An agent can prepare the brief before a call that is in the calendar. It must not start a call or send a message. It must not use data from outside the request record.

06 Practise answering objections without applying pressure

1

Fictional objection or refusal

2

Approved factual role-play

3

Human assessor reviews

4

Unaided follow-up exercise

5

Accuracy and restraint check

A training specialist writes a fictional conversation about a fee, a restriction, a misunderstanding or a decision not to buy. The assistant plays the prospective client inside the approved facts. A human assessor reviews how the salesperson explains the facts, and whether the salesperson is willing to stop.

The scenarios are approved and name no product. They come with current factual material and a marking guide. Where information is missing, the gap stays open. The trainee then practises how to refer a question, and does not invent reassurance.

Pilot with comparable scenarios, and then an exercise without the assistant. Include a clear refusal. Score factual accuracy and the disclosure of limitations, and check that the trainee escalates correctly. Stop the pilot if the simulation rewards repeated pressure or a false deadline. Stop the pilot if it rewards a claim that the risk is gone.

The sales training lead is accountable. Assess whether the trainee explains the facts accurately and respects the decision of the person, which can be a decision not to buy. Simulation feedback does not certify professional competence, and it gives no one authority to provide regulated advice.

First working session

Practise a conversation in which the person declines. Assess whether the salesperson respects that decision.

06 Build notes: objection practice

Starter prompt

"You play a prospective client in a practice conversation. Use only the attached scenario and facts. Raise the objection in the scenario in your own words. If the salesperson gives a fact that is not in the approved facts, do not accept it as true. If the scenario says that you decline, decline clearly and do not change your decision. After the role-play, list where the salesperson was accurate and where the salesperson disclosed a limitation. List also each time the salesperson applied pressure after your refusal."

Tools

Role-play: Claude or Microsoft 365 Copilot, in text or in voice mode.

Scenarios: approved fictional cases and marking guides in a restricted SharePoint folder.

Assessment: the marking guide in Microsoft Forms, completed by a human assessor.

Skills

Staff skill: explain a limitation plainly and accept a refusal.

Agent skill: a prospect-role skill with the scenario format and the rule "a refusal stays a refusal".

Method

Deliberate practice with feedback from a human assessor. In the Kirkpatrick model, the exercise without the assistant is the Level 3 (behaviour) test.

Agent limits

Level 0. Each session is a practice exchange that a trainee starts. No agent contacts a real prospect. A simulation score must not go into a sales ranking or a competence record.

07 Prepare a marketing-content review checklist

Draft marketing copy

Approved product facts

Flag claims and omissions

Authorised reviewer decides

Internal review checklist

The assistant reviews draft copy against the approved product facts and the current review checklist of the firm. It flags unsupported claims, missing qualifications, expired terms and places where a headline disagrees with its supporting text. Then it passes those passages to a reviewer.

The assistant needs the complete draft, the intended audience and channel, the current product documents and the approval requirements. Unsupported claims stay flagged. The assistant must not replace an unsupported claim with a plausible benefit.

Pilot on fictional advertisements that have known omissions and misleading comparisons. Review the text together with the prominence of material qualifications. Count the missed defects and the false approval claims. Stop if the assistant labels copy as compliant or ready to publish without the authorised review.

The marketing compliance reviewer is accountable. This internal pre-check gives no regulatory clearance. The applicable rules for submission and approval stay in force, and so do the record-keeping rules. [2]

First working session

Put an unsupported claim and a hidden qualification into a fictional advertisement.

07 Build notes: content pre-check

Starter prompt

"You pre-check draft marketing copy for a human reviewer. Use only the attached draft and the approved reference documents. Make a table with four columns: passage, problem type, product fact or checklist item, and question for the reviewer. Use these problem types only: Unsupported claim, Missing qualification, Expired term, and Headline differs from text. Say where a qualification is less prominent than the claim. Do not rewrite the copy. Do not replace a claim with a different benefit. Do not say that the copy is compliant or approved."

Tools

Draft and facts: the marketing draft and the approved product library, current versions.

Pre-check: Claude or Microsoft 365 Copilot, with the reference documents attached.

Approval: the existing marketing approval workflow, with the record of the authorised reviewer.

Skills

Staff skill: trace each claim to a product fact.

Agent skill: a content-pre-check skill with the four problem types and the rule "flag, do not rewrite".

Method

Claim-by-claim substantiation: each claim needs a source and, where one applies, a qualification with equal prominence. Test with seeded defects and count the misses.

Agent limits

Level 1. An agent can pre-check each new draft and prepare the table for the reviewer. It must not approve or publish copy. It must not submit material to a regulator.

08 Prepare a consented newsletter audience for review

Proposed newsletter audience

Check consent and suppression

Permitted and not suppressed?

Audience for human review

Hold contact out

Yes

No

The assistant compares a proposed newsletter audience with the recorded topic preferences and the purpose-specific permission. It also checks the suppression lists. Its output is a review list that gives the reason why each person is in or out. It sends nothing.

Use the approved subscriber records, the permission evidence, the purpose of the newsletter and the current suppressions. Do not infer consent from a business card or an account opening. An earlier service enquiry is not consent. When the evidence conflicts or is missing, hold the contact out until an authorised person resolves it.

Test on synthetic records that have an opt-out and a changed preference, with a duplicate address. Count the inappropriate inclusions, and check that suppressions hold across the applicable lists. Stop the pilot if the list includes a contact who withdrew permission. Stop the pilot if someone uses confidential information to infer a promotional topic.

The marketing data owner is accountable. A person approves the audience and the content separately, and approves the send process as a third step. Check the suppressions again immediately before any authorised send. The list gives no permission for an unrelated campaign. [4]

First working session

Test a synthetic audience that has an opt-out and a permission conflict, with a duplicate address.

08 Build notes: newsletter audience

Starter prompt

"You prepare a newsletter audience for human review. Use only the attached extracts. For each contact, write In or Out and give the reason. A contact is In only if permission for this purpose exists and no suppression applies. If the evidence conflicts or is missing, write 'Hold' and say what is missing. Treat duplicate addresses as one contact. Do not infer consent. Do not use data about accounts or services to select a topic. Do not send a message."

Tools

Extracts: subscriber and suppression exports from the email platform, for example Mailchimp or HubSpot, and permission evidence from the CRM.

Match: Excel or Power BI, with Claude or Microsoft 365 Copilot to explain each decision.

Review list: a SharePoint list for the marketing data owner.

Skills

Staff skill: read permission evidence for one purpose.

Agent skill: an audience-check skill with the status rules and the hold rule.

Method

Suppression-first list building: apply opt-outs before any other rule. Test with synthetic records, and check the suppressions again immediately before a send.

Agent limits

Level 1. An agent can prepare the review list for each proposed audience. It must not send a newsletter or add a subscriber. A send is a separate, approved step with its own check.

09 Learn from recorded reasons an opportunity did not proceed

Defined opportunity group

De-identified stated reasons

Separate reasons from guesses

Two analysts review

Process question to test

An analyst works from de-identified records and groups the reasons that clients gave. Examples are poor timing and an unclear condition. A request that the firm could not meet is a third example. The assistant keeps a restricted source reference, and it separates what the client said from what an employee assumed.

First, define the opportunity group, the reporting period, the result definitions and the permitted reason codes. "No response" and "unknown reason" must stay separate categories. Record silence as silence. Do not read it as price resistance or lost trust.

Two analysts check the classifications on a reviewed sample. Measure their disagreement and the reasons that no one can support. Report how many records have no reliable reason. Stop if the model identifies a client from a rare case. Stop if anyone treats the sample as representative of the full market.

The sales improvement lead is accountable. Use the findings to select a question or a process to test. The findings cannot show that a new script will recover a particular amount of revenue.

First working session

Classify a reviewed sample that has stated reasons, with records of no response and of unknown reason.

09 Build notes: reasons not to proceed

Starter prompt

"You are a sales analyst. The attached records are de-identified. Give each record one of the permitted reason codes. Use 'Stated' only when the record has the words of the client. Use 'Assumed' when an employee gave the reason. Use 'No response' and 'Unknown' as separate codes. Quote the passage that supports each code. Report how many records have no reliable reason. Do not guess a reason from silence. Do not try to identify a client. Do not estimate lost revenue."

Tools

De-identification: Microsoft Presidio, with a manual check by the analyst.

Classification: Claude or Microsoft 365 Copilot, with the de-identified extract only.

Agreement check: a spreadsheet with the codes of the two analysts.

Skills

Staff skill: separate what the client said from what an employee assumed.

Agent skill: a reason-code skill with the permitted codes and the silence rule.

Method

Content analysis with a fixed code list, and Cohen's kappa for the agreement between the two analysts. Report the share of records with no reliable reason.

Agent limits

Level 1. An agent can code the de-identified extract and prepare the summary for the analysts. It must not read records that carry client identifiers. It must not contact a person who did not proceed.

10 Check the handover from sales to the next authorised team

Enquiry and disclosures

Evidence and next decision

Receiving team checks

Accepted transfer pack

At the handover from sales, the assistant builds a handover pack. The pack has the enquiry, the verified product information that sales supplied, the disclosures that sales gave and the evidence that is still missing. It names the next decision and the team that has authority to make it. The person who receives the pack decides if it is sufficient.

Use the approved opportunity record, the source versions, the permission evidence and the checklist of the new team. Leave out personal data that no one needs. Missing evidence for identity or advice stays visibly unresolved in the pack, and so does missing application evidence. No text from the assistant can replace that evidence.

Test on fictional transfers to an adviser or an application-processing team. Include one case that must stop. Count the returned packs and the repeated requests, and count the unsupported claims that the client agreed. Stop if a handover implies suitability or acceptance that no one established. Stop if it implies a transaction instruction.

The sales workflow ends at this handover, and the sales-to-service transition lead is accountable for it. Advice and onboarding each have their own approvals and records, and so does continuing service.

First working session

Ask the new team to check one fictional handover that has required evidence missing.

10 Build notes: sales handover

Starter prompt

"You prepare a handover pack from sales to the next team. Use only the attached opportunity record and the checklist of the new team. Use four headings: Enquiry, Product information supplied, Disclosures given and Evidence still missing. Give the source and the version for each item. Name the next decision and the team that has authority to make it. If required evidence is missing, put 'NOT READY' at the top. Do not say that a product is suitable. Do not say that the client agreed unless the record shows it. Do not include personal data that the new team does not need."

Tools

Record: the opportunity record in the CRM, with source versions.

Pack: Microsoft 365 Copilot in Word, or Claude with the record and the checklist attached.

Acceptance: the intake checklist of the new team, as a form or a SharePoint list.

Skills

Staff skill: check a pack against the intake checklist of the new team.

Agent skill: a sales-handover skill with the four headings and the "not ready" rule.

Method

A definition of ready that the new team owns: the pack moves only when every required item is present. Count the returned packs.

Agent limits

Level 1. An agent can prepare the pack when a salesperson starts a handover. It must not move the opportunity to the new team. It must not start an application or an onboarding step.

Choose the first pilot

These screening scores assume approved documents and fictional opportunities, with named reviewers. They assume no integration that sends messages or changes production records. The scores are our first assessments, and we did not test them in any company. Score them again for the real workflow.

Case | Value | Feasible | Data | Effort | Risk

01 Enquiry routing | 3 | 4 | 4 | Small | Moderate

02 Product comparison | 4 | 4 | 4 | Medium | Higher

03 Proposal assembly | 4 | 4 | 4 | Medium | Higher

04 Pipeline evidence | 4 | 4 | 3 | Small | Moderate

05 Call preparation | 4 | 4 | 3 | Small | Higher

06 Objection practice | 3 | 4 | 4 | Medium | Moderate

07 Content pre-check | 5 | 3 | 3 | Medium | Higher

08 Newsletter audience | 4 | 3 | 3 | Medium | Higher

09 Reasons not to proceed | 3 | 4 | 3 | Medium | Moderate

10 Sales handover | 4 | 4 | 3 | Medium | Higher

How to read the table

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

Small effort is a bounded review of documents or metadata by one process owner. Medium effort includes specialist calculations, reconciliation of permission lists, audience tests or more than one review function. Risk describes the possible consequences of an error, and it does not describe the probability.

Use cases 02 and 03 are the first choices, together with use case 04. They use approved product documents and fictional opportunities. Audience preparation waits for verified permission and suppression controls. A failure of disclosure or privacy stops a use case, whatever its score. A failure of the advice boundary does the same.

Name an owner and a measure

Case | Accountable role | Pilot measures | Agent level

01 | Sales operations lead | Incorrect routes; unnecessary data requests | 1

02 | Product owner | Unsupported facts; omitted conditions | 1

03 | Sales proposal manager | Missing documents; inconsistent figures | 1

04 | Sales operations manager | Unsupported stages; false duplicate flags | 1

05 | Sales team lead | Unsupported claims; missed questions | 1

06 | Sales training lead | Accuracy; respect for refusal; escalation | 0

07 | Marketing compliance reviewer | Missed defects; false approval claims | 1

08 | Marketing data owner | Inappropriate inclusions; suppression failures | 1

09 | Sales improvement lead | Unsupported reasons; reviewer disagreement | 1

10 | Sales-to-service transition lead | Returned packs; unsupported agreement claims | 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 offer and its evidence

Keep a dated library of approved descriptions, terms, disclosures and channel restrictions, and connect each draft to its source version. Missing or superseded material is a matter for review, and no one can improvise to fill the gap. A person must check the calculations and the meaning of each comparison.

Keep product permission, client eligibility, suitability and marketing approval as four separate questions. A source document can answer one question and leave the others open. Keep incentives and conversion targets out of the decision about how the firm reviews errors.

Protect the person behind the opportunity

For every pilot dataset, approve the purpose, the fields, the retention and the access. Before real records go anywhere, confirm the provider terms for storage, subcontractors, cross-border arrangements and model training. Test that one team cannot get the restricted opportunities of a different team. The data protection review must cover the generated profiles and notes as well as the inputs. [4]

No person and no assistant can infer wealth, distress, disability or investment tolerance from a name or a language, or from a neighbourhood. Offer a route to a person that is not digital. A person who declines must not then get more persistent automated persuasion.

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.

Keep sending and changing records separate from drafting

Start read-only. Treat instructions inside external messages or product documents as untrusted content. Any later authorised action needs recipient and permission checks and an approval record, with a way to catch duplicates. Reconcile an uncertain send result before a second attempt. The assistant must never change terms or submit applications, and it must never execute transactions.

6 From pilot to practice

Days 1 to 15

Select one internal task, such as a comparison or a proposal. A pipeline check is a third choice. Measure the errors and omissions of today, and the review time. Approve the source versions, the fictional test cases, the reviewer and the person who has stop authority. Confirm which product scope and regulatory scope apply. Set the agent level, and write the first version of the prompt.

Days 16 to 45

Run the assistant in parallel with the existing process, and check every output. Include an out-of-date offer, a missing fee, an ambiguous enquiry, an opt-out and a question that needs an adviser. Record each error and its consequence, because a count of drafts alone says little. Give a version number to each prompt and each agent skill.

Days 46 to 75

Compare equivalent work, and count all the effort for checks and corrections. Ask if the new team or the intended reader understood the result. Look at exclusions and complaints, and at cancellations, before you draw a conclusion from a commercial change. Correct the process before you add data or channels.

Days 76 to 90

The sponsor decides to continue or to stop, and a decision to continue can include revisions. The record names who maintains the sources and who monitors the work, and it gives the rollback procedure. External messages and new claims each need separate authority, and so do changes to production records. A successful internal test does not change this rule.

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.

Product family and legal entity:

Process owner:

Client question or internal sales problem:

Approved sources and their versions:

Approved calculations:

Permitted data and purpose:

Permitted channel:

Actions the assistant can perform:

Decisions reserved for authorised people:

Agent level and tool permissions:

Prompt version and agent skill version:

Baseline and start group:

Comparison period:

Quality threshold and reviewer:

Stop condition and escalation route:

Decision date and evidence required:

Measure value without invented savings

Record the time for preparation, checks, correction and administration on comparable tasks. Count the unsupported claims, the missing conditions, the inappropriate audience inclusions and the rejected handovers. Each rate needs its own denominator. Enquiries and proposals are different groups, and completed applications are a third group. Do not compare a rate from one group with a rate from a different group.

Conversion is a commercial measure. Alone, it does not show comprehension or suitability, and it does not show causation. Log the changes in product terms, channel, season and opportunity mix next to it. Do not attribute every difference to the assistant, and do not count every decision not to proceed as a failure.

Time released is available capacity. To count it as a cash saving, you need evidence that expenditure decreased. Include the costs of software, integration, training, source maintenance and supervision. For an agent, add the usage cost of each run. Stop expansion when a check of confidentiality or permission fails, or a disclosure check fails. An improvement in average preparation time does not change this rule.

Record | What to enter

Current process | Preparation; review; correction effort; material omissions

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

Costs | Software; integration; training; source upkeep; supervision; agent usage

Sales quality | Supported claims; complete conditions; permission; accepted handovers

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

References

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

[2] Financial Services Commission, Jamaica. 2020. Market Conduct Guidelines for Securities Dealers. SR-GUID-20/01-0026, 31 January. Sections 6.5, 8 and 10. https://www.fscjamaica.org/wp-content/uploads/2020/04/SR-GUID-26-Market-Conduct-Guidelines-Amended-FINAL-FEB-7-2020.pdf

References continued

[3] Bank of Jamaica. The Banking Services (Deposit Taking Institutions) (Customer Related Matters) Code of Conduct, 2016. Article 8, gazette pages 476E–476G. https://boj.org.jm/wp-content/uploads/2019/07/Code_of_Conduct_on_Customer_Related_Matters.pdf

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

[5] 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.