Practical AI for Data Governance 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 about the records and definitions that make data usable in a financial institution, and about the controls on them. It supports the other thirteen papers in this collection. An assistant cannot give a reliable answer when no one knows the meaning or the owner of its information, or its permitted use. [1]

We wrote this paper for data owners, reporting teams, technology managers, privacy specialists and business sponsors. The proposals help to prepare evidence and find gaps. They do not authorise a model to grant access, delete records, merge customer identities or decide that processing is lawful.

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

Two reports can give different figures for the same business question, because they use different dates or definitions, or different source systems. Ask an assistant to explain the difference, and it will produce an explanation in each case. I trust the explanation only when the assistant can trace how the firm produced each number.

Data governance starts with a plain responsibility. Someone must own the definition and approve its use. That person must also correct the problem when one appears. A catalogue has value when a colleague can find the correct source in it. A policy has value when the workflow follows it, and I think that the second test is the harder one.

I would start with one report that matters. Trace its inputs and agree its definitions, and write down what is still unresolved. Then use artificial intelligence to collect the evidence, and keep the uncertainty in view.

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. Institutional sources give the context. The ten workflows, the screening scores, the prompts and the agent levels are our proposals for local test. They do not report results from a named firm.

In this paper, artificial intelligence (AI) is software that generates content or classifies information. It can also produce predictions. An assistant prepares material for a person, and an agent can also do configured steps with connected tools. Every proposal in this paper starts with read-only access.

Metadata describes a dataset: its fields, its owner, its dates and its meaning. A schema gives field names and formats. A data catalogue helps people to find sources, and a glossary records agreed business definitions. Lineage traces information from its origin, through each transformation, to its use. A data steward maintains the agreed rules. The accountable business owner approves what the rules mean and how people use them.

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

In every use case, the assistant proposes metadata and collects evidence. Authorised people approve definitions, access, sharing, retention and production changes.

Contents

Section | Page

Executive summary | 5

1 The regional business case | 6

2 What the evidence supports | 7

3 The enterprise lens | 8

Process design: BPMN swimlanes | 9

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

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

Priority and responsibility tables | 31-32

5 The foundations lens | 33

6 From pilot to practice | 34

Appendix A Pilot decision canvas | 35

Appendix B Measure value | 36

References | 37-38

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

Executive summary

Start with a source catalogue and a business glossary for one high-value process. Lineage and quality checks come later, after the source owners confirm the definitions. No pilot can start with a copy of every customer record in a model.

The main output is a decision with its evidence attached. The decision says which source to use, what a field means, which transformation produced a result and who must correct a defect. An inferred relationship stays a suggestion until the responsible owner verifies it.

Measure if colleagues can reproduce a report and find its source, and if they resolve defects with less rework. Count the review effort and the cost to keep metadata current. A catalogue that grows while its entries become incorrect or old adds work.

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 report or process that has a named business owner, and approve the metadata that the assistant can inspect. Then name the person who approves definitions and the person who approves permissions. Name also the person who approves changes. They can be different people.

Choose a bounded process

Measure the baseline

Test with review

Decide from evidence

1 The regional business case

A firm that operates across Caribbean jurisdictions can combine reports with different currencies and report periods, and different legal-entity boundaries. Inside one country also, two systems can use the same label for different measures. Establish what the data means before you join it.

The Office of the Information Commissioner of Jamaica gives standards on purpose, minimisation, accuracy, retention, security and cross-border transfers. A dataset can be technically available and still not permitted for a new model or for a provider. The responsible specialists must confirm the proposed use first. [2]

Public regional statistics give useful context. They cannot say which customer records are duplicates or which reports still use an obsolete definition. Those two questions need authorised internal evidence, and the source register of the firm is the place to start.

The diagram shows the first sequence that we propose. We left out a region-wide readiness score, because we have no regional estimate as a basis.

Select one report

Inventory inputs

Verify definitions

Record unresolved gaps

2 What the evidence supports

The 2013 risk-data principles of the Basel Committee cover governance, data architecture, consistent definitions and reconciliation. The Committee wrote them first for systemically important banks. National supervisors can extend them to other banks in proportion to the size of a bank, and to its nature and complexity. So they do not bind every Caribbean financial business automatically. We use them as a reference point. The compliance specialists of the firm confirm which requirements apply. [3]

A January 2026 newsletter of the Basel Committee reports continued difficulty with lineage and fragmented systems, and with cross-border implementation. It notes that the outputs of artificial intelligence depend on data of high quality. The newsletter gives information. It says expressly that it introduces no new supervisory guidance. [4]

Test traceability before you automate anything. An assistant can propose a field description or point to a probable transformation. The owner must then check it against the system documentation and the code, and against the outputs. A diagram does not establish that a path is verified, and a high model confidence score does not either. The owner must follow a real value along the path.

Test whether a report can be reproduced

A second analyst must be able to name the source and the definition, and the transformation between them. Where the analyst cannot, record the gap as unresolved and leave the diagram visibly incomplete.

3 The enterprise lens

Select a report behind an important decision and work back from it. Record each input, transformation, reconciliation and approval. Mark each item as documented or inferred, and mark what stays unknown. Keep unresolved gaps visible in the diagram.

Inspect metadata first. Catalogue tasks and glossary tasks usually need field names and definitions, with sample structures and no personal data. Where examples are truly necessary, the team must use approved synthetic samples or minimised samples.

The use cases in this paper support the governance team. They must not replace the systems that already enforce permissions, retention, validation and change control. Keep proposed corrections apart from approved production changes. The prompts in the build notes are start points that each firm must test.

Request and source

Assistant output

Reviewer correction

Accepted result

A source can be technically accessible and still incorrect for the intended report. Before you approve it, record its business definition and its permitted use, with any unresolved gaps.

Confirm catalogue ownership before acceptance

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

Data catalogue lead

AI drafting assistant

Source owners

Yes

No

Catalogue review starts

Provide authorised metadata and owner records

Draft purpose, fields and source descriptions

Confirm owner, purpose and sensitivity

Entry approved?

Approve entry and its review date

Catalogue review complete

Ask sponsor to resolve ownership

Catalogue entry pending

Save the owner-approved catalogue entry

Data source catalogue review

The assistant needs no customer rows. Access to the dataset still goes through the normal approval route.

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 Build a source catalogue people can use

Approved schemas and inventories

Draft owner and source entries

Source owners confirm

Usable data catalogue

In this use case, the assistant reads approved system inventories and schemas, with their documentation, and drafts catalogue entries. Each entry records the business purpose, the owner, the fields, the update frequency, the sensitivity and the location of the authoritative source.

The assistant needs authorised metadata exports and system documentation, with owner records. It needs no customer rows. When nobody can confirm who owns a source, the assistant writes a question for the sponsor. If it assigns the source to the nearest team, the entry looks complete and is incorrect.

For the pilot, take the sources behind one recurrent report. Ask a colleague who does not know the process to find the correct dataset and explain its use. Count the missing owners and the incorrect descriptions, and record the search time. Stop if the catalogue shows a restricted system or treats an inferred purpose as approved.

The data catalogue lead is accountable. Source owners approve each entry and its review date. Access to the data itself still goes through the usual request.

First working session

Select one report and list its documented sources. Include any source that has no confirmed owner.

01 Build notes: source catalogue

Starter prompt

"You draft catalogue entries for data sources. Use only the attached metadata exports and system documentation, with the owner records. For each source, give the business purpose, the owner, the fields, the update frequency and the sensitivity. Give the location of the authoritative source. Quote the document that supports each item. Write 'Owner not confirmed' where the owner records do not name one, and write a question for the sponsor. Label a purpose that you inferred as 'Inferred: owner must confirm'. Do not assign an owner. Do not ask for customer rows. Do not describe a restricted system in more detail than the documents give. Do not mark an entry as approved."

Tools

Metadata: schema and inventory exports from the source systems, with no row data.

Catalogue: Microsoft Purview, or a SharePoint list with one entry for each source.

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

Skills

Staff skill: confirm the owner and the purpose of a source before the entry goes into the catalogue.

Agent skill: a catalogue-entry skill with the entry fields and the "Owner not confirmed" rule.

Method

A metadata-first catalogue with a findability test. A colleague who does not know the process must find the correct dataset and explain its use.

Agent limits

Level 1. An agent can draft entries from approved metadata exports, and source owners approve each entry. The agent reads no customer rows. It must not write to the approved catalogue.

02 Resolve conflicting business definitions

Definition in report A

Definition in report B

Compare exclusions and periods

Business owner resolves

Dated agreed definition

Reports can use one term and count different things. An example is "active account". Others are "completed transaction" and "customer revenue". The assistant puts the competing definitions next to each other, with their periods and exclusions, and their owners. It does not select one.

Report specifications and existing glossaries are sufficient, with the calculation documentation. Add aggregate examples that show where the figures differ. A field with a similar name can mean something completely different. If a specification is missing, the disagreement stays open until its owner supplies one.

Pilot on five disputed terms. Owners agree one definition, or they document why the firm needs more than one. Assess the work by the conflicts that stay open, and by a test: two analysts must reproduce one result. Stop if a definition changes a regulatory or contractual meaning without the responsible review.

The business data owner is accountable. The owner decides between the definitions and records the date when the selected one takes effect.

First working session

Compare two definitions that conflict. Identify the calculation evidence that can settle the difference.

02 Build notes: business glossary

Starter prompt

"You compare definitions of one business term. Use only the attached report specifications and glossary entries, with the calculation documents. For the term [term], give each definition in the words of its source, with the report, the owner, the period and the exclusions. Show in a table where the definitions differ. Show the aggregate example for each definition, if the documents give one. List the evidence that the owner needs to decide. Write 'Specification missing' where a report has none. Do not select a definition. Do not write a new definition. Do not assume that two fields with similar names mean the same thing. Do not change a regulatory or contract meaning."

Tools

Sources: report specifications and calculation documents from the report teams.

Glossary: Microsoft Purview or a SharePoint list, with a version and an effective date for each term.

Comparison: Claude or Microsoft 365 Copilot for the table, and Excel for the aggregate examples.

Skills

Staff skill: reproduce a figure from its definition, and find the exclusion that makes two reports differ.

Agent skill: a definition-comparison skill with the comparison table and the "Specification missing" rule.

Method

A side-by-side definition table for each disputed term. Two analysts then reproduce one result from the agreed definition.

Agent limits

Level 0. A data steward starts each comparison for a disputed term, and the business owner decides. The assistant must not change the glossary or a report.

03 Trace a report back to its sources

Report specification

Code and interface evidence

Trace verified source links

Engineering lead confirms

Versioned lineage map

The assistant drafts a lineage map from approved transformation code and interface documentation, with the report specifications. It connects each link to evidence and marks unverified steps clearly.

Give it authorised code extracts, schemas, job definitions and report logic, without credentials and production secrets. A design document for a planned integration is not evidence that the job runs.

Pilot on one report that has a known reconciliation. An engineer checks every connection and traces a sample value along the path. Check manual adjustments explicitly. Count the unsupported connections, and record the time to explain a discrepancy. Stop if the tool invents a source or drops an adjustment.

The data engineering lead owns the verified map. The lead keeps it under change control and records when a system change makes it out of date.

First working session

Trace one approved sample value from its source to the report. Include a manual adjustment.

03 Build notes: report lineage

Starter prompt

"You draft a lineage map for one report. Use only the attached code extracts, the job definitions, the interface documents and the report specification. List each step from the source to the report, with the source system and the target, and the transformation between them. Quote the code or the document that supports each link. Label each link: Evidence in code, Evidence in document, Planned only or Not verified. List every manual adjustment as its own step. List steps that the documents do not explain as questions for the engineer. Do not invent a source or a join. Do not treat a design document as proof that a job runs. Do not include credentials."

Tools

Evidence: approved code extracts and job definitions, for example SQL, with no credentials.

Map: Claude for the draft list of links, and Visio or Microsoft Purview for the reviewed map.

Change control: the existing change process, which records when the map goes out of date.

Skills

Staff skill: follow one real value from the source to the report, through each adjustment.

Agent skill: a lineage-draft skill with the four link labels and the rule on manual adjustments.

Method

A lineage map with evidence for each link, then a trace of one sample value along the path. Only the value trace verifies the path.

Agent limits

Level 1. An agent can draft the map from approved code extracts when a report changes, and an engineer verifies each link. The agent has no production credentials. It must not change code or a job.

04 Draft testable data-quality rules

Approved business definitions

Valid and invalid test examples

Draft rules and exceptions

Quality manager approves

Testable quality specification

Start from approved business rules. The assistant changes them into proposed checks on allowed values, required fields, date relationships and reconciliation conditions. It gives the business reason for each rule, and it lists the exceptions that a person must examine.

The inputs are signed definitions and validation requirements, with synthetic examples of valid records and invalid records. Do not derive an exclusion rule from rarity alone. Where no document gives a requirement or an exception, ask its owner. Leave the threshold blank until the owner answers.

Run the pilot in a test environment that has known defects and legitimate edge cases. A data analyst reviews the false alarms and the missed defects, and the effect of any proposed threshold. Stop if a rule removes valid records or changes data without approval.

The data quality manager is accountable. The assistant drafts tests. Approved checks run in established software that records the results. Quarantine and deletion are separate steps.

First working session

Write a valid example and an invalid example before you draft the proposed quality check.

04 Build notes: quality rules

Starter prompt

"You draft data-quality checks from approved business rules. Use only the attached definitions and validation requirements, with the synthetic examples. Write each check as one testable statement. Use four types: Allowed values, Required fields, Date relationships and Reconciliation conditions. For each check, give the business reason and the definition that it comes from. Give one valid example and one invalid example. List exceptions that a person must examine. Write 'Threshold: ask owner' where the documents give no threshold. Do not derive a rule from the fact that a value is rare. Do not propose to delete or change a record. Do not ask for production data."

Tools

Rules: the approved definitions and the validation requirements, in the glossary of the firm.

Drafts: Claude or Microsoft 365 Copilot, with synthetic examples in Excel.

Execution: the established data-quality software of the firm, or SQL under change control, in a test environment.

Skills

Staff skill: write a valid example and an invalid example before you accept a rule.

Agent skill: a quality-rule skill with the four rule types and the threshold rule.

Method

Four rule types, each with a business reason. Test each rule on seeded defects and legitimate edge cases, and report false alarms and missed defects separately.

Agent limits

Level 0. A data analyst starts each draft from an approved definition. Approved checks run in established software. The assistant must not change or delete data, and it must not quarantine a record.

05 Map records to the approved retention schedule

Catalogue record categories

Approved retention schedule

Find rules and hold exceptions

Records and legal review

Retention-review list

An assistant compares catalogue entries with the approved retention schedule of the firm. It lists the datasets that have no applicable rule or no owner, and those that have no disposal process. Where no rule applies, the list says so and the period stays blank.

The assistant uses the current schedule, the record categories, the metadata of the legal-hold register and the system descriptions. Legal and privacy specialists decide which rules and exceptions apply.

One record category with a known schedule is sufficient for a pilot. Check that legal holds and conflicting obligations stay visible. Count the categories without a match and the incorrect rule mappings. Stop if the assistant proposes deletion while a hold applies. Stop if it treats a generic rule from the internet as the authority of the firm.

The records manager is accountable. The assistant deletes nothing, and no deletion is automatic. Approved disposal goes through the existing process of authorisation and verification, with its logs.

First working session

Select one record category. Confirm the current retention schedule and any legal hold.

05 Build notes: retention mapping

Starter prompt

"You compare catalogue entries with the approved retention schedule. Use only the attached catalogue entries and schedule, with the legal-hold register. For each dataset, give the record category and the schedule rule that applies, with its reference. Write 'No rule found' where the schedule has none, and leave the period blank. Write 'Hold applies' where the register shows a legal hold. List datasets that have no owner or no disposal process. Show obligations that conflict, and do not resolve them. Do not propose a retention period. Do not use a rule from outside the schedule. Do not propose a deletion. Do not say that a dataset is ready for disposal."

Tools

Schedule: the current approved retention schedule and the record categories.

Holds: the metadata of the legal-hold register, with no case content.

Review list: a SharePoint list for the records manager and the legal reviewer.

Draft: Claude or Microsoft 365 Copilot, with the catalogue entries attached.

Skills

Staff skill: check for a legal hold before any discussion of disposal.

Agent skill: a retention-map skill with the "No rule found" and "Hold applies" labels.

Method

A mapping table from each record category to its schedule rule and its hold status. Legal and privacy specialists review every gap and every conflict.

Agent limits

Level 1. An agent can prepare the review list when the catalogue or the schedule changes. It reads metadata only. It must not delete a record or start a disposal.

06 Locate sensitive fields without widening access

Approved sensitivity categories

Permitted metadata and samples

Proposed field classifications

Privacy and security confirm

Restricted classification queue

A controlled classifier proposes which fields can contain personal or confidential information. It uses schema descriptions and approved samples. Its output is a restricted review queue, and uncertain classifications go into the queue also.

Give it authorised field metadata and the approved sensitivity categories, with a small number of carefully limited examples. Credentials and account identifiers never go to a public model, and private documents do not either.

Use a synthetic dataset that has known sensitive fields and ordinary fields. Report missed sensitive fields apart from unnecessary labels. A miss can expose a customer, and an extra label only costs review time. Stop if the scan itself exports protected content outside its approved boundary.

The privacy and security data steward is accountable. People confirm each classification and each handling rule. Only they can declare that a dataset is anonymous, whatever a classifier reports.

First working session

Label a synthetic set of fields with the privacy reviewer before you test the classifier.

06 Build notes: sensitive fields

Starter prompt

"You propose sensitivity classifications for data fields. Use only the attached field metadata and the approved sensitivity categories, with the approved samples. For each field, give one category from the approved list, and give the reason from the field description or the sample. Give a confidence label: High, Medium, Low or Not sure. Put every 'Low' and 'Not sure' field in the review queue also. Do not copy sample values into your output. Do not say that a field or a dataset is anonymous. Do not ask for more rows. Do not send any content outside this workspace."

Tools

Classifier: Presidio for personal-data patterns, in the controlled environment of the firm.

Labels: the approved sensitivity categories, for example as Microsoft Purview sensitivity labels.

Queue: a restricted SharePoint list for the privacy and security data steward.

Reasons: Claude or Microsoft 365 Copilot in the same controlled environment, on metadata only.

Skills

Staff skill: confirm a classification from the field definition and its use.

Agent skill: a field-classification skill with the approved categories and the four confidence labels.

Method

Test on a labelled synthetic set. Report missed sensitive fields and unnecessary labels as separate counts, because a miss costs more than an extra label.

Agent limits

Level 1. An agent can propose classifications from approved metadata and put them in the restricted queue. A steward confirms each one. The agent must not change access or apply a label to production data.

07 Prepare access requests for an authorised decision

Purpose and access request

Check policy and evidence

Request complete?

Authorised approver review

Request missing information

Yes

No

An assistant can help to check that an access request is complete. It looks for five items: the business purpose, the required dataset, the role, the duration and the approval evidence. It flags what is missing and compares the request with the approved access policy.

The minimum data is the request form, the role catalogue, the dataset classification and the current approval rules. The requested dataset itself stays closed.

Test on historical requests with known decisions, and remove personal details where possible. Measure how frequently the assistant finds missing evidence. Measure also how frequently it incorrectly claims that approval exists. Stop if it grants access or accepts an informal message from a manager as authority. Stop if it widens the requested scope.

The access governance owner is accountable. The designated approver decides. The existing identity system enforces the permission and its expiry.

First working session

Check one redacted access request for its purpose and approval evidence, and for its expiry.

07 Build notes: access evidence

Starter prompt

"You check that one access request is complete. Use only the attached request form, the role catalogue, the dataset classification and the approval rules. Check five items: the business purpose, the dataset, the role, the duration and the approval evidence. For each item, write 'Present' or 'Missing', and quote the request. Compare the request with the approval rule for that dataset classification, and say which approver the rule requires. Flag a request with no end date. Do not treat an informal message as approval. Do not suggest a wider scope. Do not say that the request is approved. Do not open the dataset."

Tools

Requests: the access request form, for example in ServiceNow.

Rules: the role catalogue and the approval rules for each dataset classification.

Check: Claude or Microsoft 365 Copilot, with the form and the rules attached.

Enforcement: the identity system of the firm, for example Entra ID, after the approver decides.

Skills

Staff skill: check the purpose and the end date of a request before you approve it.

Agent skill: an access-request skill with the five items and the rule on informal messages.

Method

A completeness checklist of five items, on the principle of least privilege. Test on historical requests with known decisions.

Agent limits

Level 1. An agent can prepare a completeness check for each new request, and the designated approver decides. The agent has no permission in the identity system. It must not grant or widen access.

08 Assemble the evidence for a data-sharing review

Proposed sharing purpose

Fields, terms and destinations

Verified controls and open gaps

Legal and privacy review

Data-sharing decision pack

For a proposed data-sharing arrangement, the assistant assembles one structured record. The record covers the purpose, the fields, the recipients, the locations, the retention, the security and the contract evidence. Then the assistant lists the questions that are still open for the legal and privacy reviewers, and for the security reviewer.

The assistant needs the proposed agreement, the approved field list, the provider terms and the transfer details. Unsupported assurances go in one column, and verified controls go in another.

Pilot on a completed review whose evidence is still available. Compare the assessment of the specialist with what the assistant missed and what it asserted without support. Stop if the assistant treats the existence of a contract as proof that the transfer is lawful or approved.

The data protection lead is accountable. The pack is an input to the review. Any transfer still needs the appropriate internal approval. Provider claims need evidence before anyone relies on them.

First working session

Build the evidence checklist again for one completed review. Mark the questions that its contract left open.

08 Build notes: sharing review

Starter prompt

"You assemble the evidence for a data-sharing review. Use only the attached draft agreement, the approved field list, the provider terms and the transfer details. Make one record with seven headings: Purpose, Fields, Recipients, Locations, Retention, Security and Contract evidence. Under each heading, put each statement in one of two columns: 'Verified control', with its evidence, or 'Assurance without evidence'. List the fields in the agreement that are not on the approved field list. List the questions that are still open for each reviewer. Do not say that the transfer is lawful or approved. Do not treat a signed contract as proof of a control. Do not rate the provider."

Tools

Documents: the draft agreement and the provider terms, in the contract store of the firm, for example iManage or SharePoint.

Evidence table: Excel, with the two columns for each heading.

Pack: Claude or Microsoft 365 Copilot in a restricted workspace, with Word for the reviewers.

Skills

Staff skill: ask for the evidence behind each assurance of a provider.

Agent skill: a sharing-evidence skill with the seven headings and the two columns.

Method

An evidence table with two columns: verified controls and assurances without evidence. Compare the pack with the assessment of the specialist on a completed review.

Agent limits

Level 0. The data protection lead starts each pack for one proposed arrangement, and reviewers decide. The assistant must not approve a transfer or send data to a provider.

09 Investigate possible duplicate master records

Candidate master record A

Candidate master record B

Evidence for and against match

Data steward decides

Proposed controlled correction

The existing match rules flag the possible duplicates. The assistant helps a steward to inspect them. It summarises the evidence for and against each match and names the fields that still need confirmation.

Use approved identifiers, source-system references, match criteria and a limited set of comparison fields. Similar names or addresses can belong to different people, so alone they prove nothing.

Test on synthetic pairs and on an authorised test set with labels. Count false matches and missed duplicates separately. We gave this use case the lowest data score in the table. Stop if the assistant merges records or guesses missing identifiers. Stop if it shows the data of one client to another.

The master data steward is accountable. A person approves any correction through the controlled process. Keep a history that someone can reverse, and keep the original source references.

First working session

Compare two synthetic records that have similar names but different verified identifiers.

09 Build notes: duplicate investigation

Starter prompt

"You help a data steward to inspect one possible duplicate pair. The match rules of the firm already flagged the pair. Use only the attached comparison fields for the two records. Put each field of record A next to the same field of record B. Give one result for each field: Same, Different, Missing in A or Missing in B. List the evidence for a match and the evidence against it as separate lists. Name the fields that the steward must confirm at the source. Do not treat a similar name or address as proof. Do not guess a missing identifier. Do not say that the records are the same person. Do not propose a merge."

Tools

Pairs: the output of the existing match rules in the master data tool of the firm.

Comparison: Claude or Microsoft 365 Copilot in a restricted workspace, on the limited comparison fields.

Correction: the controlled correction process, which keeps a reversible history.

Skills

Staff skill: confirm a match on verified identifiers, and never on a name alone.

Agent skill: a pair-evidence skill with the four field results and the two evidence lists.

Method

A field-by-field comparison with separate lists of evidence for and against. Test on labelled pairs, and count false matches and missed duplicates separately.

Agent limits

Level 0. A master data steward starts each inspection on one flagged pair. The assistant sees only the limited comparison fields. It must not merge or change a record.

10 Assess whether a knowledge collection is ready for an assistant

Proposed knowledge collection

Purpose and access requirements

Flag versions and ownership gaps

Knowledge owner reviews

Purpose-specific readiness

Before the firm approves a document collection for a particular assistant, it needs an inventory. In this preparation task, the assistant lists the authorised documents. It flags obsolete versions, conflicting instructions, missing owners, weak access labels and documents without a review date. Then it writes a readiness report for the business owner.

The minimum data is the approved documents, the version records, the permissions and the questions that the future assistant must answer. Add a list of topics that it must refuse.

Put known outdated documents and restricted files into the pilot, with questions that the collection cannot answer. Check that the report finds each defect. Check also that a sample answer goes back to the correct source. Stop if someone declares readiness while access conflicts or version conflicts are unresolved.

The knowledge collection owner approves readiness only for the tested purpose and audience. A collection that answers internal policy questions well can be unsuitable for client advice or model training. Uses that no one tested stay unapproved.

First working session

Test a collection that contains a deliberately obsolete version and a restricted test document.

10 Build notes: knowledge readiness

Starter prompt

"You assess a document collection before an assistant uses it. Use only the attached document inventory, the version records, the permissions and the list of intended questions. For each document, check four things: a current version, a named owner, an access label and a review date. List each document that fails a check. List pairs of documents that give conflicting instructions, and quote the two passages. For each intended question, say if the collection has a current source that answers it. List the topics that the future assistant must refuse. Write a readiness report for the stated purpose and audience only. Do not say that the collection is ready for another purpose. Do not resolve a conflict between documents. Do not change a permission."

Tools

Collection: a SharePoint library with version history and access labels, and a named owner for each document.

Inventory: Claude or Microsoft 365 Copilot, with the document list and the version records attached.

Test set: a fixed list of intended questions and refusal topics, in Excel.

Skills

Staff skill: retire an obsolete version before the assistant can read it.

Agent skill: a readiness-check skill with the four document checks and the question test.

Method

A document audit with four checks for each document, then a fixed question test. Each sample answer must go back to a current source, and each refusal topic must get a refusal.

Agent limits

Level 1. An agent can prepare the inventory and the readiness report for one collection and one purpose. The knowledge collection owner approves. The agent must not change a document or a permission.

Choose the first pilot

These screening scores assume one bounded report, approved metadata or synthetic examples, named owners and no production integration. The scores are our first assessments, and no company measured them. Score each use case again against the actual source inventory.

Case | Value | Feasible | Data | Effort | Risk

01 Source catalogue | 4 | 4 | 3 | Small | Moderate

02 Business glossary | 4 | 4 | 3 | Medium | Higher

03 Report lineage | 5 | 3 | 3 | Medium | Higher

04 Quality rules | 4 | 4 | 3 | Medium | Higher

05 Retention mapping | 4 | 3 | 3 | Medium | Higher

06 Sensitive fields | 5 | 3 | 3 | Medium | Higher

07 Access evidence | 4 | 4 | 3 | Medium | Higher

08 Sharing review | 5 | 3 | 3 | Medium | Higher

09 Duplicate investigation | 4 | 2 | 2 | Medium | Higher

10 Knowledge readiness | 4 | 4 | 3 | Medium | Higher

How to read the table

Value is the probable decrease in rework or the gain in traceability. Feasible shows how well the firm can run a controlled test. Data shows how ready the evidence is. Each score is from 1 (weak) to 5 (strong).

Small effort is a desk test on metadata only, inside one ownership area. Medium effort adds code or sample preparation and rule tests, or review across specialist functions. Risk rates the consequence of an error, and it does not rate the probability.

Use cases 01 and 02 are the first choices, for one report. Add use cases 03 and 04 after owners confirm the definitions. Use cases that touch personal data or permissions need the relevant specialist review, and so do use cases that touch disposal. No score, even a high one, overrides a failed control.

Name an owner and a measure

Case | Accountable role | Pilot measures | Agent level

01 | Data catalogue lead | Missing owners; incorrect descriptions; search time | 1

02 | Business data owner | Unresolved conflicts; reproducible results | 0

03 | Data engineering lead | Unsupported connections; omitted adjustments | 1

04 | Data quality manager | Missed defects; false alarms | 0

05 | Records manager | Unmatched categories; incorrect rule mappings | 1

06 | Privacy and security data steward | Missed sensitive fields; unnecessary labels | 1

07 | Access governance owner | Missing evidence; false approval claims | 1

08 | Data protection lead | Omitted questions; unsupported assurances | 0

09 | Master data steward | False matches; missed duplicates | 0

10 | Knowledge collection owner | Obsolete versions; unresolved access gaps | 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

Keep proposed metadata separate from approved facts

Store the suggestion of the assistant and its supporting evidence as separate records, and store the decision of the reviewer as a third. The authoritative catalogue and glossary change only when someone approves a change. Each approved change must have a version and an effective date.

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.

Minimise the data used to improve the data

Schemas and documentation come first. When the work requires row-level examples, approve the fields, the sample size, the purpose and the retention. Confirm the storage terms and the model-training terms of the provider. Use an authorised test account to confirm that the system refuses access where it must.

The data-protection review must cover the original records and the new outputs. The outputs can contain sensitive descriptions or inferred links. A catalogue of restricted sources can itself need protection. [2]

Test the boundary as well as the answer

Put a document that tells the assistant to disclose data into the test set. Add an obsolete policy and a request to change access. The assistant must inspect them as test content. It must not follow their instructions or change access. Check that logs contain no secrets. Check also that a reviewer can investigate a bad output without access to unrelated records.

Use the existing change-control process and the existing incident process. The business owner must know how to withdraw an incorrect definition and repair the reports that depend on it. The owner must also know how to tell the people that the error affects.

6 From pilot to practice

Days 1 to 15

Select one report and its accountable owner. Make an inventory of the source metadata, and note the gaps. Record how long it takes to locate a source and explain a measure, and how long it takes to resolve a defect. The later decision depends on that baseline. Agree the review boundaries and the access boundaries. Set the agent level, and write the first version of the prompt.

Days 16 to 45

Draft the catalogue and the glossary in a separate review space, and check each proposed entry against evidence. Put the difficult material in deliberately: conflicting terms and undocumented spreadsheets, with obsolete versions. Unresolved questions stay on the list. Give a version number to each prompt and each agent skill.

Days 46 to 75

Trace the report, and test the approved quality rules on non-production data. Then a second analyst reproduces the result. Count the corrections and the missed defects, and record the maintenance effort. Repair any incorrect definitions before you add systems.

Days 76 to 90

The sponsor reviews the evidence and decides to expand or not. The sponsor names the continuing steward and the source-review schedule, with the triggers for change. Production integration needs its own approval. So do any access changes or data changes that a correct draft catalogue seems to justify.

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.

Report or process and owner:

Decision the data supports:

Approved metadata and sample fields:

Definitions requiring agreement:

Actions reserved for authorised people:

Agent level and tool permissions:

Prompt version and agent skill version:

Baseline and test dataset:

Quality threshold and reviewer:

Stop condition and correction route:

Decision date and evidence required:

Measure value without inventing savings

Record the time to find and check information, and the time to correct it. Do this before the pilot and in the pilot. Compare equivalent reports, and include catalogue maintenance. Record if a second analyst can reproduce the result. Record also if the process finds the important defects.

Do not treat a count of catalogue entries as coverage. Establish the relevant source inventory first. Then report which sources have verified owners and definitions, and review dates. Show the percentage that is still unknown.

Released time is capacity, and a new use of that time does not produce a cash saving alone. Report an actual decrease in expenditure apart from operational benefits. Examples of such benefits are fewer returned reports and a faster resolution of a documented discrepancy. Keep the costs of licences, integration, supervision and maintenance in the record. For an agent, add the usage cost of each run.

Record | What to enter

Current process | Source-finding time; correction effort; reproducibility

Pilot process | Same measures on comparable reports; all review and upkeep

Costs | Licence; integration; supervision; metadata maintenance; agent usage

Governance quality | Verified owners and definitions; detected defects; unresolved gaps

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

References continued

[3] Basel Committee on Banking Supervision. 2013. Principles for effective risk data aggregation and risk reporting. January; scope, paragraphs 13–19 (pages 3–4), and Principles 1–3 (pages 6–9). https://www.bis.org/publications/201301-guidelines-principles-effective-risk-data-aggregation-and-risk-reporting.pdf

[4] Basel Committee on Banking Supervision. 2026. Implementation of the Principles for effective risk data aggregation and risk reporting. Newsletter 36, 6 January. Informational publication. https://www.bis.org/publicati

ons/implementation-principles-effective-risk-data-aggregation-and-risk-reporting-bcbs-239-principles

[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.