INVARIANT ENGINEERING

Synthetic Auditability Sample Packet

AI-assisted commercial credit memorandum drafting and policy-exception support

Six-Month Reconstruction Test v1.0

Packet version 1.0 | Methodology source reviewed July 21, 2026
Packet reviewed July 24, 2026 | Synthetic transaction SCM-2026-0042

Example Regional Bank and Example Borrower LLC are fictional.
This is not a client case study, audit opinion, legal opinion, certification, or compliance conclusion.

Contents and navigation

Use the links below to move through the reference packet. Page numbers are shown for the final rendered layout.

Document navigation

Jump to Sources • Jump to final CTAs

How to read the identifiers

EV- identifies an evidence element. GAP- identifies an observed auditability gap. REM- identifies a remediation item. SYN- and SCM- identify synthetic systems, records, and events.

Executive summary and scope

Synthetic-data notice

Everything concerning the bank, borrower, transaction, systems, records, actors, amounts, ratios, policies, timestamps, model, prompts, outputs, approvals, findings, and remediation actions is fictional and synthetic. The sample includes no real people, institutions, addresses, account numbers, customer data, personal-guarantor information, or protected-class information. The provider identity is intentionally omitted from the public sample; a synthetic internal vendor code is retained.

Purpose

This packet demonstrates the five substantive outputs of an Audit Grill-Me Sprint for one AI-assisted workflow. It shows how a consequential workflow is mapped, reconstructed from retained evidence, evaluated across eight evidence elements, summarized for senior oversight, and translated into a 30/60/90-day remediation sequence.

Method and illustrative result

The Six-Month Reconstruction Test v1.0 evaluates eight evidence elements: model version; prompt or system instruction; input data; retrieved sources; original output; human reviewer; approval or override record; and retained evidence. [1]

Only fully Present elements count toward the public evidence-readiness bands: 0–3 High reconstruction risk, 4–6 Partial auditability, and 7–8 Stronger evidence posture. [2]

Fully Present	Partial	Missing	Public band	Confidence
5	2	1	Partial auditability	Moderate

The result is an illustrative evidence-readiness result. It is not an audit rating, risk score, legal opinion, compliance conclusion, certification, or assurance opinion.

Engagement context

The public Audit Grill-Me Sprint is a $9,500 fixed-price, ten-business-day engagement covering one AI workflow and five deliverables. No sensitive production data is required for the initial sprint. The first sprint is designed to avoid sensitive production-data access wherever possible. [3]

The five deliverables represented here are: AI Workflow Evidence Map; Six-Month Reconstruction Test; Auditability Gap Register; Audit Committee Briefing Memo; and 30/60/90-Day Remediation Backlog.

Professional boundary

Professional boundary

Invariant Engineering provides technical AI auditability engineering, evidence-readiness assessments, and remediation support. We do not provide legal advice, issue independent audit opinions, certify compliance, or act as the client’s external auditor.

Workflow boundary

The assessed workflow is AI-assisted commercial credit memorandum drafting and policy-exception support. The AI may assemble supplied financial-analysis inputs, retrieve identified policy sections, draft specified memorandum language, and identify possible policy exceptions for analyst disposition.

The AI must not approve or deny credit; set pricing, terms, or risk ratings; accept or reject a policy exception; produce an adverse-action reason; communicate a credit decision; initiate loan booking; or replace analyst or credit-officer judgment.

Regulatory and risk-management context

Regulation B is cited only to explain why a commercial-credit workflow may be audit-sensitive. This packet makes no determination about applicability, triggered duties, or compliance. CFPB Circular 2022-03 is not used as current guidance because the CFPB lists it as withdrawn on May 12, 2025. [4] [5]

Federal Reserve SR 26-2 supersedes SR 11-7. The sample uses selected governance concepts only as context and makes no conclusion about applicability. NIST AI RMF 1.0 and NIST AI 600-1 are voluntary risk-management context, not compliance criteria or scoring inputs. [6] [7] [8]

Artifact 1 — AI Workflow Evidence Map

Stable HTML anchor: #ai-workflow-evidence-map

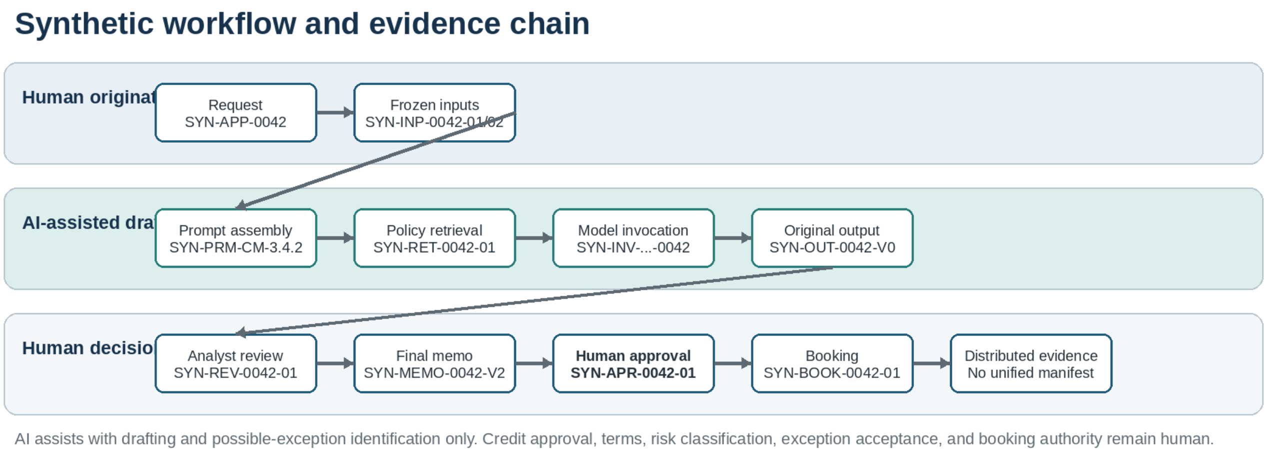

Figure 1. Synthetic workflow and evidence chain.

Evidence-map conclusion

Strong evidence exists for the model deployment, frozen inputs, original AI output, analyst review and edit lineage, and final human approval. Partial evidence exists for the exact rendered prompt and precise retrieved policy context. The retained records do not form a complete, governed, time-bounded evidence chain, and no authoritative manifest binds all eight elements to the transaction.

Roles and decision authority

Actor	Fictional role	Workflow responsibility	Authority
ACT-01	Relationship Manager A	Initiates request and narrative context	No final credit authority
ACT-02	Credit Analyst A	Freezes inputs, invokes MemoAssist, reviews and edits, confirms possible exceptions	May recommend; may not finally approve
ACT-03	Credit Officer A	Reviews final memo, accepts exceptions, imposes conditions, approves	Final synthetic approval authority
ACT-04	Loan Operations Specialist A	Performs non-AI closing checks and booking	Cannot alter the credit decision

Systems and evidence stores

System	Name	Primary evidence	Recovered at six months	Owner
SYN-SYS-01	Commercial Loan Origination System	Request, exceptions, approval	Yes	OWN-01
SYN-SYS-02	Financial Spreading Workbench	Structured financial snapshot	Yes	OWN-01
SYN-SYS-03	MemoAssist Application	Run, prompt selection, review, edits	Yes, except rendered prompt	OWN-03
SYN-SYS-04	Credit Policy Library	Versioned policy documents	Yes	OWN-04
SYN-SYS-05	Policy Retrieval Index	Retrieval metadata	Partial	OWN-02
SYN-SYS-06	Hosted Model Gateway	Invocation configuration and original response	Yes	OWN-02

System	Name	Primary evidence	Recovered at six months	Owner
SYN-SYS-07	Credit File Repository	Final memo and approval packet	Yes	OWN-01
SYN-SYS-08	Core Lending Booking System	Booked-facility record	Yes	OWN-01

Retention for every synthetic system: Not established in the synthetic evidence set. Provider identity: Intentionally omitted from the public sample; synthetic internal vendor code retained.

Model, prompt, and retrieval configuration

Field	Synthetic value
Application	SYN-APP-CM-2.8.1
Model record	SYN-MOD-01
Provider code	SYN-TP-LLM-01
Model revision	SYN-GEN-A-2025-11-17
Deployment	SYN-DEP-CREDIT-MEMO-R3
Gateway config	SYN-CFG-2026-01-03
Temperature / top-p	0.2 / 0.9
Maximum output	3,500 tokens
Prompt template	SYN-PRM-CM-3.4.2
Corpus	SYN-CORP-CREDIT-2026-01-05.1
Invocation	SYN-INV-20260112-100518-0042
Retrieval event	SYN-RET-0042-01

Evidence-element summary

ID	Element	Status	Observed evidence	Gaps
EV-01	Model version	Present	Invocation, provider code, revision, deployment, configuration	GAP-06
EV-02	Prompt / system instruction	Partial	Template retained; exact rendered prompt absent	GAP-01, GAP-07
EV-03	Input data	Present	Frozen structured snapshot and narrative	GAP-07
EV-04	Retrieved sources	Partial	Corpus and policy versions retained; chunk trace absent	GAP-02
EV-05	Original output	Present	Complete unedited model response, timestamp, and invocation link were produced.	GAP-03
EV-06	Human reviewer	Present	Reviewer identity, displayed version, review timestamp, edit lineage	GAP-05, GAP-03
EV-07	Approval / override record	Present	Approval, exceptions, conditions, basis, timestamp	GAP-05, GAP-03
EV-08	Retained evidence	Missing	Individual records exist; integrated governed chain does not	GAP-03, GAP-04

Detailed map

The complete 16-stage workflow-stage evidence map is preserved in Appendix A, split across smaller landscape tables to avoid unreadable word wrapping.

Artifact 2 — Six-Month Reconstruction Test

Stable HTML anchor: #six-month-reconstruction-test

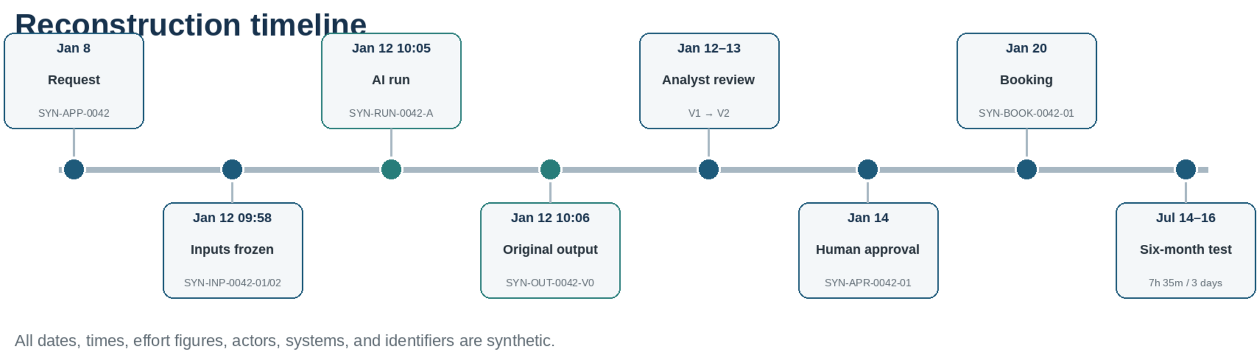

Figure 2. Synthetic transaction and reconstruction timeline.

Test objective

Reconstruction question

Six months after the event, can Example Regional Bank reconstruct how the AI-assisted commercial credit memorandum for synthetic transaction SCM-2026-0042 was produced, reviewed, changed, approved, and connected to the downstream action using retained evidence rather than interviews, memory, or recreated context?

Sampled task and outcome

Object	Synthetic record or formulation
Human business question	Whether to approve the requested facility renewal and increase under human-originated terms
AI task	Draft specified memorandum sections and identify possible policy exceptions
Original AI artifact	SYN-OUT-0042-V0
Final analyst artifact	SYN-MEMO-0042-V2
Human approval	SYN-APR-0042-01
Downstream action	SYN-BOOK-0042-01
Reconstruction request	SYN-REC-0042-01
Manual evidence index	SYN-EVIDX-0042-01

Credit Officer A approved the synthetic $8.5 million facility for a synthetic 24-month term, accepted two analyst-confirmed policy exceptions, and imposed monthly borrowing-base certificates and quarterly customer-concentration reporting. The AI did not approve the facility or set pricing, terms, or risk classification.

Original and edited artifacts

Original AI output — SYN-OUT-0042-V0

Example Borrower LLC (fictional) requests renewal and increase of its revolving facility from $6.0 million (synthetic) to $8.5 million (synthetic) for 24 months (synthetic). Adjusted debt to EBITDA is 3.61x (synthetic), above the 3.50x (synthetic) guideline. Largest-customer concentration is 28% (synthetic), above the 25% (synthetic) threshold. Both exceptions require Credit Committee approval. No material litigation was identified.

Final analyst memorandum — SYN-MEMO-0042-V2

Example Borrower LLC (fictional) requests renewal and increase of its revolving facility from $6.0 million (synthetic) to $8.5 million (synthetic) for 24 months (synthetic). Adjusted debt to EBITDA is 3.61x (synthetic), above the 3.50x (synthetic) guideline. Largest-customer concentration is 31% (synthetic), based on the dated receivables concentration report, above the 25% (synthetic) threshold. The leverage and concentration items are analyst-confirmed policy exceptions. Analyst recommendation: if approved by an authorized credit officer, condition the facility on monthly borrowing-base certificates and quarterly customer-concentration reporting. All pricing, terms, risk classification, exception acceptance, and approval decisions are human-originated and human-approved.

Reconstructed edit lineage

Edit	Original condition	Analyst treatment	Evidence
Customer concentration	28% older narrative value	Corrected to 31% dated source value	Reviewer diff and input snapshot
Approval authority	AI stated both exceptions required Credit Committee approval	Replaced with policy-supported delegated-authority formulation	Reviewer diff and policy citations
Litigation assertion	Unsupported statement included	Removed	Reviewer diff
Possible exceptions	AI identification only	Analyst independently confirmed both	Final memo and exception records
Monitoring conditions	Not included	Added by analyst and adopted by credit officer	V2 and approval record
Decision ownership	Could be read as authoritative	Reserved decision fields to human personnel	V2

Recoverable and unrecoverable evidence

Category	Reconstruction result
Produced	Request; frozen financial inputs; application and prompt-template versions; model deployment and invocation; corpus and policy versions; original output; analyst identity and edits; exception records; approval; monitoring conditions; booking.
Not produced	Exact rendered prompt; exact field ordering; proof of how the 28% and 31% values were presented; exact retrieved chunks and text; chunk IDs; ranking; scores; presentation order; authoritative evidence manifest.
Not established	Retention periods, hold treatment, authorized-deletion procedures, and periodic producibility testing.

Reconstruction effort

Activity	Synthetic effort
Locate origination and final credit-file records	55 minutes
Export and verify input snapshots	1 hour 10 minutes
Retrieve invocation and original output	1 hour 5 minutes
Correlate review and edit records	1 hour 20 minutes
Recover policy versions and corpus metadata	1 hour 25 minutes
Resolve identifiers and assumptions	55 minutes
Assemble manual evidence index	45 minutes

Total: 7 hours 35 minutes over three synthetic business days.

Eight-element result

ID	Element	Result	Status
EV-01	Model version	Invocation, provider code, revision, deployment, configuration	Present
EV-02	Prompt / system instruction	Template retained; exact rendered prompt absent	Partial
EV-03	Input data	Frozen structured snapshot and narrative	Present
EV-04	Retrieved sources	Corpus and policy versions retained; chunk trace absent	Partial
EV-05	Original output	Complete unedited model response, timestamp, and invocation link were produced.	Present
EV-06	Human reviewer	Reviewer identity, displayed version, review timestamp, edit lineage	Present
EV-07	Approval / override record	Approval, exceptions, conditions, basis, timestamp	Present
EV-08	Retained evidence	Individual records exist; integrated governed chain does not	Missing

Overall conclusion

Partial auditability — Moderate confidence

The human-approved outcome, final memorandum, original AI output, analyst edits, exceptions, and approval event can be reconstructed. The precise AI context cannot be reconstructed in full because the rendered prompt and exact retrieved policy chunks were not retained. The organization also cannot demonstrate a complete, governed, time-bounded evidence chain. Locating the final memorandum is not equivalent to reconstructing how the AI-assisted artifact was produced.

Artifact 3 — Auditability Gap Register

Stable HTML anchor: #auditability-gap-register

Auditability gap summary

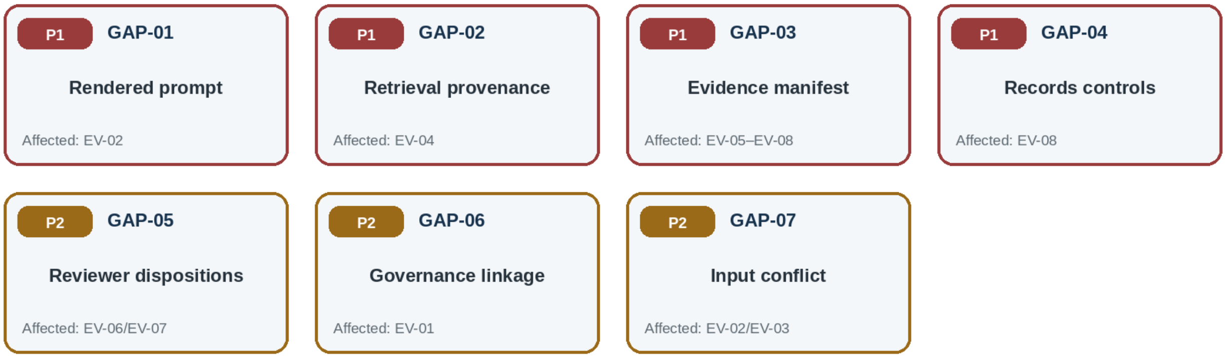

Four reconstruction-critical gaps (P1) and three material hardening gaps (P2). No numerical risk score is assigned.

Figure 3. Gap population and priority summary.

Prioritization method

Priority	Definition
P1 — Reconstruction-critical	Blocks or materially undermines reconstruction, has repeat workflow exposure, and lacks strong compensating evidence
P2 — Material hardening gap	Degrades reliability, governance, or reviewer proof but has meaningful compensating evidence
P3 — Improvement opportunity	Improves efficiency or clarity without materially changing the present reconstruction conclusion

No numerical risk score is calculated. No legal or compliance conclusion is assigned to any gap.

GAP-01 — Exact rendered prompt not retained

Field	Detail
Evidence / priority	EV-02 | P1 / High
Affected stage	Prompt assembly and invocation
Observed condition	The exact fully rendered prompt was not retained. Precise field ordering and presentation of the conflicting concentration values cannot be established.
Compensating evidence	Template, system instruction, input snapshots, application version, original output
Owner / remediation	OWN-02; OWN-03 | REM-01 | First 30 days
Validation	Capture the exact rendered prompt and reproduce it without source code, interviews, or memory.
Residual limitation	Prompt capture does not establish output accuracy.

GAP-02 — Retrieval provenance incomplete

Field	Detail
Evidence / priority	EV-04 | P1 / High
Affected stage	Policy retrieval
Observed condition	Corpus and policy-document versions are known, but exact chunks, text, rank, scores, and presentation order are absent.
Compensating evidence	Versioned policies, corpus ID, retrieval event, output citations
Owner / remediation	OWN-02; OWN-04 | REM-02 | Days 31–60
Validation	Reproduce exactly what policy context the model received.
Residual limitation	Provenance does not prove retrieval sufficiency or interpretation.

GAP-03 — No authoritative transaction evidence manifest

Field	Detail
Evidence / priority	EV-05–EV-08 | P1 / High
Affected stage	Cross-system linkage
Observed condition	No authoritative manifest binds invocation, prompt, inputs, retrieval, output, review, approval, and downstream action.
Compensating evidence	Common transaction ID, timestamps, versions, manual index
Owner / remediation	OWN-03; OWN-01 | REM-03 | Days 31–60
Validation	Reconcile an automatically generated manifest to all source-system records.
Residual limitation	A manifest depends on source-record completeness and integrity.

GAP-04 — Retention and records controls not established

Field	Detail
Evidence / priority	EV-08 | P1 / High
Affected stage	Evidence retention and producibility
Observed condition	The synthetic evidence set does not establish approved classification, retention, hold, authorized-deletion, or periodic producibility controls.
Compensating evidence	Individual records were available at six months
Owner / remediation	OWN-05 | REM-04 | Days 61–90
Validation	Test a retained synthetic record under organization-approved rules, including hold and authorized disposition.
Residual limitation	Records controls cannot recover evidence that was never captured.

GAP-05 — Reviewer dispositions are not structured

Field	Detail
Evidence / priority	EV-06; EV-07 | P2 / Moderate
Affected stage	Analyst review
Observed condition	The review captures identity and edits but does not require accept, reject, or modify dispositions for each material AI item.
Compensating evidence	Edit history, final memo, exception records, approval basis
Owner / remediation	OWN-01; OWN-03 | REM-05 | First 30 days
Validation	Inject an unsupported statement and block submission until a disposition and rationale are recorded.
Residual limitation	Structured dispositions do not prove the quality of judgment.

GAP-06 — Invocation not linked to governance records

Field	Detail
Evidence / priority	EV-01 | P2 / Moderate
Affected stage	Invocation and change governance
Observed condition	The deployment is identifiable but not directly linked to AI inventory and deployment-change approval records.
Compensating evidence	Model revision, deployment, gateway config, application version, timestamp
Owner / remediation	OWN-06; OWN-02 | REM-06 | Days 31–60
Validation	Trace a selected invocation directly to inventory, permitted use, owner, and change record.
Residual limitation	Governance linkage does not establish output accuracy.

GAP-07 — Conflicting inputs not resolved before generation

Field	Detail
Evidence / priority	EV-02; EV-03 | P2 / Moderate
Affected stage	Input preparation
Observed condition	MemoAssist accepted conflicting 28% and 31% values without requiring resolution or an explicit warning.
Compensating evidence	Analyst corrected the conflict before approval
Owner / remediation	OWN-03; OWN-01 | REM-07 | First 30 days
Validation	Repeat the conflict and verify that silent precedence is impossible.
Residual limitation	Conflict controls still depend on source-authority rules or human review.

Gap-to-remediation traceability

Evidence	Status	Gap IDs	Remediation IDs
EV-01	Present	GAP-06	REM-06
EV-02	Partial	GAP-01, GAP-07	REM-01, REM-07
EV-03	Present	GAP-07	REM-07
EV-04	Partial	GAP-02	REM-02
EV-05	Present	GAP-03	REM-03
EV-06	Present	GAP-05, GAP-03	REM-05, REM-03
EV-07	Present	GAP-05, GAP-03	REM-05, REM-03
EV-08	Missing	GAP-03, GAP-04	REM-03, REM-04, REM-08

Artifact 4 — Audit Committee Briefing Memo

Stable HTML anchor: #audit-committee-briefing-memo

To	Audit Committee of Example Regional Bank
From	Synthetic AI Auditability Review
Subject	Reconstruction of AI-assisted commercial credit memorandum SCM-2026-0042
Date	July 24, 2026

Executive conclusion

The sampled AI-assisted commercial credit memorandum could be reconstructed in outline but not in full. Five of eight evidence elements were fully Present, two were Partial, and one was Missing. The public band is Partial auditability, with Moderate overall reconstruction confidence.

The final human-approved outcome is well supported. The exact rendered prompt, exact retrieved policy context, unified evidence manifest, and organization-approved records controls were not produced or established.

Why the workflow matters

The AI-generated draft was not the credit decision, but it became an input to a consequential human process. The original output used the older 28% customer-concentration value, misstated approval authority, and included an unsupported litigation assertion. The analyst corrected or removed each item before approval. Human review operated as an important control; the auditability question is whether that control and the preceding AI context can be reconstructed later.

Principal strengths

- Model invocation and deployment identity were recoverable.
- Both frozen input records were preserved.
- The complete unedited original AI output was produced.
- Analyst identity, reviewed version, timestamps, and edit lineage were traceable.
- Final human approval, exceptions, conditions, basis, and approved artifact version were documented.

Principal gaps

- GAP-01: the exact rendered prompt was not retained.
- GAP-02: exact retrieved chunks, ranking, scores, and order were unavailable.
- GAP-03: evidence had to be assembled manually across eight systems.
- GAP-04: approved retention, hold, authorized-disposition, and producibility controls were not established.
- GAP-05 through GAP-07: reviewer dispositions, governance linkage, and input-conflict handling require hardening.

Oversight implications

Management should distinguish runtime evidence from governance documentation; a final human artifact from the original AI artifact; a retained template from the exact rendered prompt; document references from precise retrieved context; and surviving records from an organization-approved retention and evidence-production design.

Immediate management actions

1. REM-01 — Capture the exact rendered prompt.
2. REM-05 — Require structured reviewer dispositions.
3. REM-07 — Prevent silent precedence when inputs conflict.
4. Begin planning REM-03 — Create an authoritative evidence manifest.
5. Begin planning REM-04 — Establish organization-approved records controls.

Questions for the committee

1. Which owner is accountable for producing the complete evidence chain for one AI-assisted transaction?
2. Can management demonstrate the exact prompt and retrieved context for a sampled output?
3. Does the reviewer workflow document why each material AI suggestion was accepted, rejected, or modified?
4. Can runtime invocations be traced directly to inventory and deployment-change records?
5. Has the organization approved how long each evidence element must be retained and how it will be produced?
6. When will management retest the workflow from retained evidence without interviews, memory, or source-code reconstruction?

Limitations and professional boundary

This review covers one fictional workflow and one synthetic transaction. It does not assess credit prudence, validate financial statements, evaluate model performance, fairness, bias, cybersecurity, or third-party security, determine legal or supervisory applicability, assess compliance, certify the workflow, or provide an independent audit opinion.

Professional boundary

Invariant Engineering provides technical AI auditability engineering, evidence-readiness assessments, and remediation support. We do not provide legal advice, issue independent audit opinions, certify compliance, or act as the client’s external auditor.

Artifact 5 — 30/60/90-Day Remediation Backlog

Stable HTML anchor: #30-60-90-day-remediation-backlog

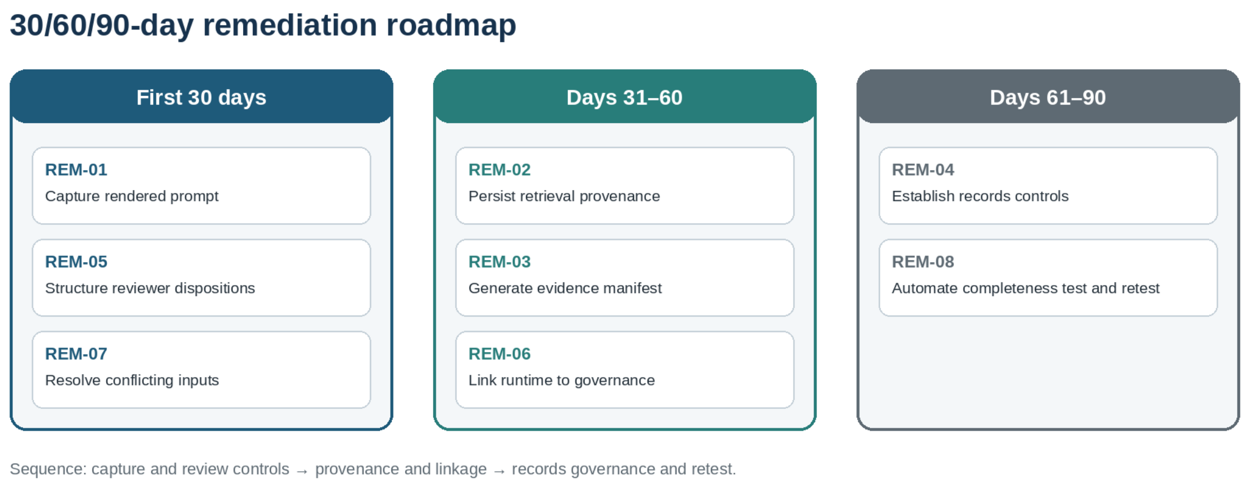

Figure 4. Remediation sequence.

Backlog objective

The sequence moves the workflow from manually reconstructed records toward a more complete, linked, governed, and reproducible evidence trail. It prioritizes direct evidence capture and review controls before broader records governance and automated retesting.

REM-01 — Persist the exact rendered system and user prompt, assembled field values, template and application versions, model-configuration reference, transaction ID, and content digest

Field	Detail
Period / gap	First 30 days | GAP-01
Owner	OWN-02; OWN-03
Task	Persist the exact rendered system and user prompt, assembled field values, template and application versions, model-configuration reference, transaction ID, and content digest.
Expected output	Invocation export containing exact rendered prompt and direct transaction link.
Retest	Reproduce the exact prompt without source code, interviews, or recreated context.
Complexity	Medium
Residual limitation	Prompt capture does not establish output correctness.

REM-05 — Add structured accept, reject, or modify dispositions for every possible exception, unsupported assertion, conflicting input, and material AI suggestion

Field	Detail
Period / gap	First 30 days | GAP-05
Owner	OWN-01; OWN-03
Task	Add structured accept, reject, or modify dispositions for every possible exception, unsupported assertion, conflicting input, and material AI suggestion.
Expected output	Review record with disposition, rationale, displayed version, identity, timestamp, and attestation.
Retest	Inject an unsupported statement and block submission until disposition.
Complexity	Low to medium
Residual limitation	Structured dispositions do not prove judgment quality.

REM-07 — Detect conflicting structured and narrative values; require analyst resolution or an explicit unresolved-conflict marker before generation

Field	Detail
Period / gap	First 30 days | GAP-07
Owner	OWN-03; OWN-01
Task	Detect conflicting structured and narrative values; require analyst resolution or an explicit unresolved-conflict marker before generation.
Expected output	Conflict-resolution event showing values, authoritative source, actor, and timestamp.
Retest	Repeat the 28% / 31% conflict and verify silent precedence is impossible.
Complexity	Medium
Residual limitation	Source authority still depends on organization rules or human review.

REM-02 — Persist corpus ID, policy document and version, chunk ID, exact presented text or immutable reference, rank, score, and presentation order

Field	Detail
Period / gap	Days 31–60 | GAP-02
Owner	OWN-02; OWN-04
Task	Persist corpus ID, policy document and version, chunk ID, exact presented text or immutable reference, rank, score, and presentation order.
Expected output	Retrieval-trace export.
Retest	Reconstruct the exact policy context without rerunning retrieval.
Complexity	Medium
Residual limitation	Provenance does not prove retrieval sufficiency.

REM-03 — Generate a versioned transaction evidence manifest binding all eight elements and the approval and booking records

Field	Detail
Period / gap	Days 31–60 | GAP-03
Owner	OWN-03; OWN-01
Task	Generate a versioned transaction evidence manifest binding all eight elements and the approval and booking records.
Expected output	Manifest with evidence IDs, timestamps, source systems, versions, relationship keys, and content digests.
Retest	Reconcile the manifest to every source record and test missing or mismatched relationships.
Complexity	Medium
Residual limitation	Manifest quality depends on source records.

REM-06 — Link every invocation to the approved AI inventory record, deployment-change record, stated use limitation, and oversight owner

Field	Detail
Period / gap	Days 31–60 | GAP-06
Owner	OWN-06; OWN-02
Task	Link every invocation to the approved AI inventory record, deployment-change record, stated use limitation, and oversight owner.
Expected output	Invocation evidence containing inventory and change-record references.
Retest	Trace a controlled invocation to the authorized deployment and governance records.
Complexity	Low to medium
Residual limitation	Governance linkage does not establish output accuracy.

REM-04 — Approve and implement fictional organization-specific evidence classification, retention, hold, authorized-deletion, and producibility controls

Field	Detail
Period / gap	Days 61–90 | GAP-04
Owner	OWN-05
Task	Approve and implement fictional organization-specific evidence classification, retention, hold, authorized-deletion, and producibility controls.
Expected output	Organization-approved records-control mapping and successful configuration/test evidence.
Retest	Produce a retained synthetic record and test hold and authorized-disposition behavior.
Complexity	Medium
Residual limitation	Records controls cannot recover evidence never captured.

REM-08 — Add an automated evidence-completeness check and perform a fresh Six-Month Reconstruction Test simulation

Field	Detail
Period / gap	Days 61–90 | GAP-01–GAP-07
Owner	OWN-03; validated by OWN-07
Task	Add an automated evidence-completeness check and perform a fresh Six-Month Reconstruction Test simulation.
Expected output	Completeness report, exception report, and retest packet.
Retest	Reconstruct a fresh synthetic transaction using retained evidence only.
Complexity	Medium
Residual limitation	A successful retest applies to the tested workflow and configuration.

Remediation acceptance criteria

1. Identify the exact model deployment and its governance record.
2. Produce the exact rendered prompt.
3. Produce frozen input records.
4. Reproduce the exact retrieved policy context.
5. Produce the original unedited output.
6. Identify the reviewer and the artifact reviewed.
7. Produce structured reviewer dispositions and final approval evidence.
8. Link the downstream action to the approved artifact.
9. Produce the evidence chain under organization-approved retention and records controls.

A fresh test must report the evidence result that actually exists. Completion of remediation tasks does not itself establish a Stronger evidence posture.

Consolidated remediation schedule

Period	Remediation	Gap linkage	Owner
First 30 days	REM-01	GAP-01	OWN-02; OWN-03
First 30 days	REM-05	GAP-05	OWN-01; OWN-03
First 30 days	REM-07	GAP-07	OWN-03; OWN-01
Days 31–60	REM-02	GAP-02	OWN-02; OWN-04
Days 31–60	REM-03	GAP-03	OWN-03; OWN-01
Days 31–60	REM-06	GAP-06	OWN-06; OWN-02
Days 61–90	REM-04	GAP-04	OWN-05
Days 61–90	REM-08	GAP-01–GAP-07	OWN-03; validated by OWN-07

Sources and next steps

The following sources provide the narrow regulatory and voluntary risk-management context cited in this packet. They are not Audit Grill-Me scoring criteria and do not create a compliance conclusion.

1. Invariant Engineering — Audit Grill-Me methodology
   Eight evidence elements and evidence-first reconstruction method.
2. Invariant Engineering — Audit Grill-Me worksheet
   Public evidence-readiness scoring bands.
3. Invariant Engineering — Audit Grill-Me Sprint engagement details
   Public engagement scope, pricing, duration, and deliverables.
4. Consumer Financial Protection Bureau — Equal Credit Opportunity Act (Regulation B)
   Narrow context for why a commercial-credit workflow may be audit-sensitive.
5. Consumer Financial Protection Bureau — Withdrawn Guidance
   Official record showing CFPB Circular 2022-03 was withdrawn on May 12, 2025.
6. Federal Reserve — SR 26-2: Revised Guidance on Model Risk Management
   Contextual governance source; no applicability conclusion.
7. NIST — Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1
   Voluntary AI risk-management framework.
8. NIST — Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile, NIST AI 600-1

Voluntary cross-sectoral companion profile for generative AI.

Next steps

View engagement details →
`/engagements`

Review the methodology →
`/audit-grill-me-methodology`

Apply for a fit review →
`/apply?source=engagements&intent=audit-grill-me-sprint`

Appendix A — Detailed workflow-stage evidence map

The original 16-stage evidence map is split into four smaller landscape tables. Every field is synthetic.

Workflow stages 1–4

Stage and record	Expected and observed evidence	System / owner / linkage	Integrity and status
1. Request initiated SYN-APP-0042	Expected: Request ID, requester, amount, term, timestamp Observed: Request record with direct transaction link	SYN-SYS-01 / OWN-01 Linkage: Direct	Integrity: No material conflict Status: Complete
2. Financial inputs prepared SYN-INP-0042-01	Expected: Frozen structured snapshot, source date, ratios Observed: Snapshot with revenue, EBITDA, debt, coverage, concentration, aging	SYN-SYS-02 / OWN-01 Linkage: Direct	Integrity: No conflict within snapshot Status: Complete
3. Narrative input prepared SYN-INP-0042-02	Expected: Frozen relationship narrative and timestamp Observed: Narrative contains older 28% concentration value	SYN-SYS-03 / OWN-01 Linkage: Direct	Integrity: Conflicts with 31% dated snapshot Status: Complete as input; conflict unresolved
4. Prompt selected SYN-PRM-CM-3.4.2	Expected: Template, system instruction, app version, field references Observed: Template and instructions recovered	SYN-SYS-03 / OWN-03 Linkage: Linked to run	Integrity: Template integrity supportable Status: Complete

Workflow stages 5–8

Stage and record	Expected and observed evidence	System / owner / linkage	Integrity and status
5. Prompt rendered Runtime prompt	Expected: Exact assembled prompt and values Observed: Exact rendered prompt not retained	SYN-SYS-03 / OWN-02, OWN-03 Linkage: Partially inferable	Integrity: Field order and conflict presentation cannot be proven Status: Missing
6. Policies retrieved SYN-RET-0042-01	Expected: Corpus, documents, chunks, ranks, scores, order Observed: Corpus and policy versions recovered; exact chunks and ranking absent	SYN-SYS-05 / OWN-02, OWN-04 Linkage: Partial	Integrity: Exact model context cannot be reproduced Status: Partial
7. Model invoked SYN-INV-20260112-100518-0042	Expected: Provider code, revision, deployment, config, timestamp, link Observed: All invocation details recovered	SYN-SYS-06 / OWN-02 Linkage: Direct	Integrity: Not directly linked to inventory/change approval Status: Complete with governance weakness
8. Original output generated SYN-OUT-0042-V0	Expected: Complete unedited response, timestamp, invocation link Observed: Complete unedited response, timestamp, and invocation link were produced.	SYN-SYS-06 / OWN-02 Linkage: Direct	Integrity: No content digest is published Status: Complete

Workflow stages 9–12

Stage and record	Expected and observed evidence	System / owner / linkage	Integrity and status
9. Analyst review begins SYN-REV-0042-01	Expected: Reviewer, displayed version, timestamp, review activity Observed: Reviewer identity and displayed version recovered	SYN-SYS-03 / OWN-01, OWN-03 Linkage: Direct	Integrity: No structured disposition for every assertion Status: Complete with review weakness
10. Analyst edits output SYN-MEMO-0042-V1	Expected: Versioned edit history and supported corrections Observed: Concentration corrected; assertion removed; authority revised	SYN-SYS-03 / OWN-01 Linkage: Direct	Integrity: Edit lineage recoverable Status: Complete
11. Final memorandum submitted SYN-MEMO-0042-V2	Expected: Final analyst artifact and timestamp Observed: Final memo recovered	SYN-SYS-07 / OWN-01 Linkage: Direct	Integrity: Not a complete evidence package by itself Status: Complete
12. Exceptions recorded SYN-EXC-0042-01 / 02	Expected: Exception, policy threshold, actual, reviewer, status Observed: Both exception records recovered	SYN-SYS-01 / OWN-01, OWN-04 Linkage: Direct	Integrity: No structured AI-suggestion disposition requirement Status: Complete

Workflow stages 13–16

Stage and record	Expected and observed evidence	System / owner / linkage	Integrity and status
13. Human approval SYN-APR-0042-01	Expected: Approver, artifact version, exceptions, conditions, basis, timestamp Observed: All approval fields recovered	SYN-SYS-01 / OWN-01 Linkage: Direct to final memo	Integrity: AI evidence linked through separate systems Status: Complete
14. Closing checks ACT-04 non-AI checks	Expected: Evidence booking prerequisites completed Observed: Represented as completed; detail outside scope	SYN-SYS-08 / OWN-01 Linkage: Linked to approval	Integrity: Outside AI-assistance boundary Status: Not applicable to AI score
15. Booking SYN-BOOK-0042-01	Expected: Booked amount, term, approval link, timestamp Observed: Booking record recovered	SYN-SYS-08 / OWN-01 Linkage: Direct to transaction and approval	Integrity: End-to-end lineage requires assembly Status: Complete as downstream action
16. Evidence retention Distributed records	Expected: Governed immutable time-bounded evidence chain Observed: Individual records available; no manifest or retention mapping	Multiple / OWN-05, OWN-03 Linkage: Manual	Integrity: Long-term completeness and producibility unproven Status: Missing

Appendix B — Synthetic record and configuration detail

Synthetic identifier registry

Identifier	Type	Meaning
SCM-2026-0042	Sample transaction	Primary transaction reconstructed
SYN-APP-0042	Credit request	Renewal and increase request
SYN-INP-0042-01	Input	Frozen structured financial snapshot
SYN-INP-0042-02	Input	Frozen relationship-manager narrative
SYN-RUN-0042-A	AI run	MemoAssist orchestration run
SYN-INV-20260112-100518-0042	Invocation	Hosted-model gateway invocation
SYN-RET-0042-01	Retrieval	Policy retrieval event
SYN-OUT-0042-V0	Original output	Unedited model draft
SYN-REV-0042-01	Review	Analyst review session
SYN-MEMO-0042-V1	Analyst edit	First edited memorandum
SYN-MEMO-0042-V2	Final analyst artifact	Memorandum submitted for approval
SYN-EXC-0042-01	Exception	Adjusted-leverage exception
SYN-EXC-0042-02	Exception	Largest-customer-concentration exception
SYN-APR-0042-01	Approval	Final human approval with conditions
SYN-BOOK-0042-01	Downstream action	Booking after approval and non-AI checks
SYN-REC-0042-01	Reconstruction	Six-month request
SYN-EVIDX-0042-01	Manual index	Assembled during test; not original workflow evidence

Transaction facts and synthetic thresholds

Fact	Synthetic value	Source
Existing facility	$6.0 million	SYN-APP-0042
Requested facility	$8.5 million	SYN-APP-0042
Requested term	24 months	SYN-APP-0042
Financial period end	December 31, 2025	SYN-INP-0042-01
Revenue	$42.8 million	SYN-INP-0042-01
Adjusted EBITDA	$5.1 million	SYN-INP-0042-01
Total debt	$18.4 million	SYN-INP-0042-01
Adjusted debt / EBITDA	3.61x	SYN-INP-0042-01
Fixed-charge coverage	1.31x	SYN-INP-0042-01
Largest-customer concentration — dated source	31%	SYN-INP-0042-01
Largest-customer concentration — older narrative	28%	SYN-INP-0042-02
A/R over 90 days	4.8%	SYN-INP-0042-01
Pricing	Not stated; human-only	—
Risk rating	Not stated; human-only	—

Policy measure	Synthetic threshold	Result
Adjusted debt-to-EBITDA guideline	3.50x	Actual 3.61x; exception confirmed
Largest-customer concentration	25%	Actual 31%; exception confirmed
Fixed-charge-coverage minimum	1.20x	Actual 1.31x; no exception

Event timeline and reconstruction effort

Synthetic timestamp (ET)	Event	Record
Jan 8, 2026 9:12 a.m.	Request opened	SYN-APP-0042
Jan 9, 2026 2:20 p.m.	Financial records represented	Preliminary source records
Jan 12, 2026 9:58 a.m.	Inputs frozen	SYN-INP-0042-01/02
Jan 12, 2026 10:05:18 a.m.	AI run begins	SYN-RUN-0042-A; invocation
Jan 12, 2026 10:05:19–10:05:21 a.m.	Policy retrieval	SYN-RET-0042-01
Jan 12, 2026 10:06:02 a.m.	Original output returned	SYN-OUT-0042-V0
Jan 12, 2026 1:47:36 p.m.	Analyst corrects and revises	SYN-REV-0042-01; V1
Jan 13, 2026 9:16 a.m.	Credit officer requests clarity	Human review comment
Jan 13, 2026 11:02:19 a.m.	Final memo submitted	V2
Jan 14, 2026 3:31:08 p.m.	Human approval	SYN-APR-0042-01
Jan 20, 2026 12:10 p.m.	Booking after closing checks	SYN-BOOK-0042-01
Jul 14, 2026 8:30 a.m.	Reconstruction request opened	SYN-REC-0042-01
Jul 16, 2026 4:05 p.m.	Manual index completed	SYN-EVIDX-0042-01

Synthetic policy registry

Policy ID	Title	Effective date	Relevant rule
SYN-POL-CCP-2025.4	Commercial Credit Policy	Oct. 1, 2025	3.50x leverage guideline
SYN-POL-CONC-2025.2	Customer Concentration Addendum	Nov. 15, 2025	25% concentration threshold
SYN-PROC-MEMO-2025.3	Credit Memorandum Procedure	Dec. 1, 2025	Review and approval-authority presentation
SYN-STD-AIDRAFT-2025.1	AI-Assisted Drafting Standard	Dec. 15, 2025	Human review; no autonomous credit decisions

Locked synthetic instructions

System instruction

You draft commercial credit memorandum language from supplied records and identify possible policy exceptions. You do not approve or deny credit, set pricing or terms, assign a risk rating, create adverse-action reasons, or replace analyst or credit-officer judgment. Cite only supplied sources. Mark uncertain or conflicting values for analyst review.

User-template instruction

Draft the requested sections for synthetic transaction SCM-2026-0042 using the structured financial snapshot, relationship-manager narrative, and retrieved policy excerpts. List possible exceptions separately. Do not resolve conflicting facts; surface them.

Appendix C — Assumptions, limitations, and traceability

Assumptions

Assumption	Locked treatment
Source records	The identified synthetic source records exist exactly as described.
Timestamps	Source-system timestamps are treated as correctly associated with their records.
Actor authentication	Actor IDs are assumed to resolve to authenticated fictional roles.
Financial accuracy	The packet does not independently validate the synthetic financial statements.
Credit quality	The packet does not assess whether the human credit decision was prudent.
Model performance	No statistical accuracy, hallucination rate, bias, fairness, robustness, or benchmark assessment is performed.
Cybersecurity	No penetration test, security audit, or third-party security assessment is performed.
Legal application	No legal determination is made concerning Regulation B or any other law.
Supervisory application	No conclusion is made about whether SR 26-2 applies.
Framework conformity	No NIST AI RMF or Generative AI Profile conformity claim is made.
Provider identity	Intentionally omitted from the public sample; synthetic internal vendor code retained.
Retention	Individual records were available at six months; required periods are not established in the synthetic evidence set.
Deletion	Deletion procedures are not established in the synthetic evidence set.
Public handling	The packet contains only synthetic information; no confidentiality guarantee is made.
Assurance	The packet is not an independent assurance report.

Limitations

- The evidence-readiness result applies only to the fictional workflow and synthetic transaction described.
- The packet does not validate the borrower’s financial statements or determine whether the credit decision was prudent.
- The packet does not evaluate model accuracy, hallucination frequency, bias, fairness, robustness, benchmark performance, cybersecurity, or third-party security.
- The packet does not determine whether Regulation B or another law applies and does not conclude compliance or noncompliance.
- The packet does not determine whether SR 26-2 applies and does not claim conformity with NIST AI RMF 1.0 or NIST AI 600-1.
- The packet does not establish required retention periods, identify a real provider, provide a confidentiality guarantee, issue an audit opinion, provide legal advice, certify the workflow, or guarantee acceptance by a regulator, auditor, customer, or court.

Full evidence-gap-remediation crosswalk

Evidence	Element	Status	Record or condition	Gaps	Remediations
EV-01	Model version	Present	SYN-INV-20260112-100518-0042; SYN-MOD-01; SYN-DEP-CREDIT-MEMO-R3	GAP-06	REM-06
EV-02	Prompt / system instruction	Partial	SYN-PRM-CM-3.4.2; rendered prompt absent	GAP-01; GAP-07	REM-01; REM-07
EV-03	Input data	Present	SYN-INP-0042-01; SYN-INP-0042-02	GAP-07	REM-07

Evidence	Element	Status	Record or condition	Gaps	Remediations
EV-04	Retrieved sources	Partial	SYN-RET-0042-01; corpus and documents present; chunk trace absent	GAP-02	REM-02
EV-05	Original output	Present	SYN-OUT-0042-V0	GAP-03	REM-03
EV-06	Human reviewer	Present	SYN-REV-0042-01; V0-to-V2 lineage	GAP-05; GAP-03	REM-05; REM-03
EV-07	Approval / override	Present	SYN-APR-0042-01; SYN-EXC-0042-01/02	GAP-05; GAP-03	REM-05; REM-03
EV-08	Retained evidence	Missing	Individual records available; governed chain absent	GAP-03; GAP-04	REM-03; REM-04; REM-08

Professional boundary

Professional boundary

Invariant Engineering provides technical AI auditability engineering, evidence-readiness assessments, and remediation support. We do not provide legal advice, issue independent audit opinions, certify compliance, or act as the client’s external auditor.