DHISORA ADVISORY

AI Playbook for Accounting and Finance

Judgment and Human Oversight

Connecting judgment, authority, risk, controls and evidence in practice

A synthesized operating guide for CFOs, executives and process owners

10 October 2026 | Version 1.0.1

From principles to practice

Choose a valuable use case, design authority and controls, test before deployment, record the reasons for decisions, and prepare to stop and restart.

Includes 12 practical chapters, two hypothetical cases and eight editable templates.

DHISORA Advisory — Articles and Playbook

Publication scope and permitted use

Public Reference Edition | Version 1.0.1 | 10 October 2026

Purpose and scope

This playbook supports education, discussion, assessment and the design of AI use in accounting and finance. It provides general guidance and does not review or verify any particular entity's facts, documents or systems.

Professional scope and application

The content is not entity-specific legal, tax, investment or accounting advice, an audit report, an assurance report, system certification or confirmation of compliance with laws or standards. Before application, appropriately responsible and authorized people must review the facts, data, risks, segregation of duties and applicable requirements, obtaining relevant specialist input as the work requires.

Examples templates and outcomes

Hypothetical cases illustrate concepts and are not reports of services delivered to actual clients. Templates are starting points for design: define entity-specific criteria, data, authority and responsibilities before use. Use of this publication does not guarantee business performance, IPO approval, acceptance by auditors or regulators, or the effectiveness of controls or AI systems.

References and engagements

Check the versions in force, their applicability and official sources at the time of use. Referencing a framework or standard does not mean every provision applies to every entity, or that its issuer endorses DHISORA. Downloading, reading or an initial inquiry does not constitute an engagement or confirmation of acceptance. Scope, deliverables, fees and responsibilities must be agreed separately.

Publication scope and permitted use continued

Public Reference Edition | Version 1.0.1 | 10 October 2026

Reading citation and internal use

You may read and download the publication for reference and adapt its templates for use within your own organization. Identify the DHISORA source and version in adapted copies and clearly mark your own modifications. Quotations must preserve meaning and context and must not imply DHISORA’s endorsement of your work or conclusions. This does not limit exceptions or rights provided by law.

Redistribution and third party rights

Share the website link to the latest edition where possible. Public redistribution of the complete publication, translation or adaptation for publication, sale, inclusion in commercial products, services or courses, and use of names or marks to imply endorsement require prior written permission, except where permitted by law. Third-party standards, publications and marks remain subject to their respective owners’ rights. DHISORA’s permission does not extend to those rights.

User responsibilities and liability scope

Users must act lawfully and consistently with their duties, review information before deciding, and not use this publication to make false or misleading statements or claim an endorsement that does not exist. This notice explains the publication’s scope. It does not remove anyone’s rights or exclude duties or liability that cannot lawfully be excluded.

Edition and corrections

Version 1.0.1, published 10 October 2026, adds publication-scope and permitted-use notices and PDF formatting to the version 1.0 Thai playbook and its English translation. To report an error or request permission, identify the page, wording and version and contact dhira@dhisoraadvisory.com.

Contents

1 Goals and how to use this playbook....5
2 Shared terminology and the article series....6
3 Choose human oversight for the context....7
4 Roles authority and segregation of duties....8
5 Connect risks and controls....9
6 Implementation lifecycle and decision gates....10
7 Test and decide on readiness....11
8 Everyday operation and evidence....12
9 Stop contain and restart....13
10 Hypothetical expected credit loss estimate....14
11 Hypothetical high volume receipt matching....15
12 Measure value and expand on evidence....16
Template 1 Use Case Card....17
Template 2 Risk and Control Matrix....18
Template 3 Authority and Oversight Record....19
Template 4 Judgment Record....20
Template 5 Test Plan and Results....21
Template 6 Go Live Decision Record....22
Template 7 Incident and Restart Record....23
Template 8 Management Monitoring Review....24
Governance and control references....25
1 COSO GenAI 2026....25
2 NIST AI RMF and GenAI Profile....25
3 ISO IEC 42001 2023....25
4 EU AI Act....25
5 Bank of Thailand AI policy 12 September 2025....25
6 IIA Statements of Position....25
Professional and human oversight references....26
7 TFAC Financial Reporting Standards....26
8 TFAC Auditing Standards....26
9 European AI HLEG 2019....26
10 Australian DTA Key Terms....26
11 Federal Reserve SR 26 2....26
Editorial review and content direction....27

1 Goals and how to use this playbook

The value of AI in finance should be assessed through better decisions and operating results within appropriate boundaries, authority and controls. Output speed alone does not capture all the value, and a human signature does not guarantee that an output is correct or reliable.

This playbook is designed for CFOs, executives, process owners and those charged with governance. Use it to select a process, define roles and authority, design controls, test, deploy, monitor, stop and restart. It connects with business process services (BPS), financial closing and organizational readiness without assuming that every entity uses AI or the same reporting framework.

Table 1 Goals and how to use this playbook

Starting step	Questions to answer
Select one process	What problem needs solving, who owns the process, and what outcome can be measured?
Define the scope	What may AI propose or execute, what data may it use, and what is prohibited?
Establish a baseline	Current preparation and review time, errors, rework and total cost.
Connect risk and authority	What could go wrong, who is responsible, where are decisions made, and what system permissions apply?
Pilot and decide	Acceptance criteria, evidence, limitations and the person authorized to approve deployment.
Monitor and adapt	Alert recipients, stop triggers, fallback arrangements and restart authority.

Reading guide Chapters 2–5 establish the principles and structures; Chapters 6–9 address operational design; Chapters 10–11 present hypothetical cases; and Chapter 12 measures value. Templates 1–8 help develop entity-specific records. The final sections distinguish reference sources from the editorial review of the original material.

> Scope These are proposals for assessment and design, not a legal opinion, certificate or assurance report. Involve technical, security or legal specialists as needed, and assess professional scope and independence before accepting an engagement.

2 Shared terminology and the article series

Judgment means reasoned assessment and choice in the presence of information, uncertainty and constraints. AI may help generate options, retrieve information or analyze it, but the responsibilities of people and organizations must remain clear. Judgment need not be defined as something AI will never be able to perform.

Human oversight means providing the roles, information, time, competence and authority that enable people to oversee a process effectively. AI governance establishes direction and accountability. Internal control brings objectives and risk responses into everyday work. These concepts are connected, but they are not interchangeable.

Table 2 Shared terminology and the article series

Topic	Core question	Connected outcome
Human Oversight	Where should a person decide or intervene?	Oversight approach and the capabilities of the person overseeing it.
Delegation of Authority (DOA)	Who has authority to decide what?	Aligned decision authority, conditions and system permissions.
Risk Management	What could go wrong, and what would it affect?	Risks, responses and residual risk.
Process Design	How will this work in every reporting period?	Process steps, controls, responsible owners and evidence.
Executive Synthesis	How do we see the whole picture and choose a starting point?	Connect the four sets of principles.
This playbook	How do we start, test and operate?	Evidence-based decisions supported by working templates.

Conditions for sound judgment The decision maker needs access to information and limitations, sufficient time, competence, authority within the defined scope, and incentives to challenge appropriately. Disagreement need not be manufactured every time, but the reasons must be reviewable.

> Reading the references Distinguish requirements binding on the entity from frameworks, guidance and DHISORA proposals. The tables, templates and hypothetical cases in this book are a synthesis; they are not forms endorsed by COSO or another authority.

3 Choose human oversight for the context

Loop terminology does not have one mandatory definition across all frameworks. European ethics guidance discusses HITL, HOTL and HIC, while Australia's Digital Transformation Agency (DTA) also describes HOOTL. The definitions below explain how these terms are used for design in this playbook. [9–10]

Table 3 Choose human oversight for the context

Approach	How it operates	Evidence needed
Human-in-the-loop (HITL)	A person considers the result before a specified action is taken.	Evidence, time, competence and authority to approve or reject.
Human-on-the-loop (HOTL)	The system operates within boundaries; a person monitors and intervenes.	Alerts reach a person who can stop the process before unacceptable consequences occur.
Human-out-of-the-loop (HOOTL)	No person participates in the specified decision cycle.	Tested boundaries, a responsible owner, monitoring and an incident plan.
Human-in-command (HIC)	People determine goals, policy, authority and how the system is used.	Actual authority to change, restrict or stop use, with oversight of overall outcomes.

Choosing an approach Consider impact, speed, detectability, reversibility, the potential extent of harm and review capacity together. Work volume or value per transaction is not the only criterion. HIC is a governance layer across the other approaches, not a fourth level of automation.

> When control is not yet adequate Restrict AI to proposing results, reduce volume or permissions, change the process, or suspend use. Adding HITL without providing more time, information or competence can turn review into a rubber stamp.

A monthly retrospective review may serve another control purpose, but it does not demonstrate that HOTL can intervene in time. Test alert delivery, response and stopping times in the actual activities covered by the design.

4 Roles authority and segregation of duties

AI is not a person to whom authority is delegated. DOA defines decision authority; RACI defines working roles as Responsible, Accountable, Consulted and Informed. Signing authority and system permissions must be checked separately. Being Accountable in a RACI matrix does not automatically confer legal approval authority.

Management owns and operates controls. Risk or compliance functions provide support and challenge within their roles. Internal audit provides independent assurance. The board and audit committee exercise governance within their mandates; the audit committee should not be classified as a second-line operational function. [6]

Table 4 Roles authority and segregation of duties

Matter to define	Authority and conditions	Safeguards against conflicting roles
Approve the use case scope	The authorized decision maker understands the risks and limitations.	Separate the proposer and approver as warranted by risk.
Create or change rules and versions	Define permitted changes and changes requiring fresh approval.	A person changing a rule should not approve their own change without a safeguard.
Significant items and exceptions	Determine the decision level from substance, impact and requirements.	Clearly distinguish preparation, review and decision making.
Suspend or restrict the system	Define who can stop it, by which channel and within what scope.	Avoid approval delays that prevent stopping within the specified time.
Accept residual risk	The authorized person considers the evidence and reasons.	Distinguish this decision from a preparer's or reviewer's signature.
Restart or expand use	Decide using test results and the revised limitations.	Challenge the results and identify who will follow up.

Approving a rule does not replace individual transaction approvals reserved by law, policy or other conditions. Risk acceptance does not permit a breach of mandatory requirements. An exception does not always require escalation by exactly one level; consider the substance of the decision.

Smaller entities may combine some roles. Identify conflicts and use supplementary review or an independent reviewer as appropriate. Both the delegator and the delegate retain responsibilities within their roles; responsibility should not be pushed entirely upward or downward.

5 Connect risks and controls

Start with the process objective. Identify causes, potential failure events and their effects, then assess data, systems, people, decisions and service providers. Judgment risk is an explanatory lens for risks to decisions in this playbook, not a universally mandated risk category.

Distinguish inherent risk from the risk remaining after controls. Do not lower a risk rating merely because an approver has been named. Assess control design and operating results. Acceptable risk must remain within mandatory requirements and defined authority.

Distinguish risk appetite, the types and levels of risk an entity is willing to take to pursue its objectives, from risk tolerance, the acceptable variation in performance, and risk capacity, the ability to absorb consequences. Definitions and thresholds must align with the entity's framework. Do not set limits above its capacity or mandatory constraints.

Table 5 Connect risks and controls

COSO component	DHISORA application	What should be observable
Control Environment	Roles, competence, ethics and incentives.	Authority and resources to challenge results.
Risk Assessment	Assess the use case and changes.	Impact, correlated errors and residual risk.
Control Activities	Design control points and permissions.	Prevention, detection, containment and correction.
Information and Communication	Provide information needed for decisions.	Sources, limitations, alerts and escalation.
Monitoring Activities	Monitor and evaluate deficiencies.	Remediation, a responsible owner, due dates and retesting.

COSO's 2026 GenAI guidance applies the existing internal control framework to generative AI. This table connects concepts for design; it is not a mapping of every principle or a compliance certification. [1]

Potential extent of harm Consider throughput, detection and stopping times, correlated errors and the ability to reverse actions. Aggregate value during the period of exposure is one perspective, not a formula that replaces assessment of every risk. Financial limits do not replace constraints concerning data, fraud or affected parties' rights.

> Responses Avoidance may remove risk from a particular activity, while risks involving people, data or fallback arrangements remain. Detection can occur before or after an event; connect it to the time available to reduce actual consequences.

6 Implementation lifecycle and decision gates

This six-stage lifecycle is DHISORA's proposed working approach. Stages may overlap or be revisited. Approval at each stage must follow the entity's authority arrangements; the board need not approve every use case or every process step.

Table 6 Implementation lifecycle and decision gates

Stage	Deliverables	Decision gate
1 Define goals and scope	Problem, process owner, requirements and prohibitions.	A real need exists and falls within the engagement scope.
2 Discover and register	The current process (As-Is), actual usage and shadow AI.	Relevant data points, systems and owners are identified.
3 Assess and choose	Risk, oversight, benefits and costs.	Reasons support the chosen approach or the decision not to use AI.
4 Design and prepare	The future process (To-Be), risk and control matrix (RCM), authority, permissions and fallback.	Critical points have a responsible owner and a means of checking them.
5 Test and decide on deployment	Test results, deficiencies and residual risk.	Predefined criteria are met, or permissible conditions can be approved.
6 Operate and review	Actual evidence, metrics, incidents and changes.	Operation remains controlled within boundaries, or is restricted and reassessed.

The To-Be process should show inputs, proposals or actions, checks, decision points and handoffs. A separate AI swimlane can help explain the process but is not mandatory. Accountability remains with people or organizations according to their roles.

Changes requiring review Source data, models or providers, significant prompts or rules, permissions, volumes, activity types and requirements may change. Assess material changes and test within the relevant scope. Do not assume that testing an earlier version covers a new version.

> Approval is not required at every handoff. Select controls that address risks and link document identifiers appropriately. Relationships among use cases, risks and controls may be many-to-many.

7 Test and decide on readiness

Set acceptance criteria before test results are known, together with scope, sample size and selection method, period, tester, expected results and the rationale for adequate coverage. Distinguish design, implementation and operation during the tested period. Passing in one period does not establish effectiveness throughout the year or in every environment.

Table 7 Test and decide on readiness

Test dimension	Coverage to consider	Example evidence
Data and business logic	Complete data, wrong periods, duplicates, omissions and boundary cases.	Compare results with independently verifiable sources.
Outputs and review	Plausible but incorrect results; whether reviewers see limitations.	Corrections, reasons and independent checks.
Permissions and changes	Prohibited actions, new rules and new versions.	Permission tests and change records.
Alerts and stopping	Undelivered alerts, absent responses and failed stops.	Actual alert, response and stopping times.
Fallback and recovery	Reversal, pending items and resuming work.	Impact and completeness after switching methods.
Volume and cost	Reviewer workload, backlog and total cost.	Baseline compared with the pilot period.

Simulate errors in an approved environment separate from live accounting. Do not introduce incorrect entries into the live ledger. The previous method may also be wrong, so agreement with it is not always conclusive. Effective IT general controls (ITGC) do not establish that application-specific logic is correct.

> Decision outcomes Approve, approve with conditions, defer, or decline use. Conditions need scope, limits, an expiry date, a responsible owner and suspension triggers. Noncompliance with mandatory requirements cannot be substituted with risk acceptance.

Do not change criteria retrospectively to obtain a pass. If existing criteria prove unsuitable, record the reasons, assess the effects, obtain approval under the applicable authority, and test transparently against the new criteria.

8 Everyday operation and evidence

Evidence should arise during the work and reflect what was actually considered. Link sources, periods, methods, assumptions, material alternatives, contradictory information, checks, corrections, decision makers and relevant versions. There is no need to retain every prompt or answer without a purpose. A model's internal message history is not a substitute for verifiable reasons.

Table 8 Everyday operation and evidence

Activity	Control	Retrievable evidence
Receive data	Check sources, completeness, period and permissions.	Dataset identifiers, time, control totals and check results.
Prepare or propose	Use approved versions and rules within scope.	Proposed output, material assumptions and limitations.
Review and decide	Examine information and alternatives according to risk.	Checks performed and reasons for accepting, changing or rejecting.
Execute	Check authority and permissions for the activity.	Actual transaction, operator, time and outcome.
Handle exceptions	Assign an owner; distinguish data correction from risk acceptance.	Issue, scope, reasons and closure status.
Monitor and retain	Check alerts and the quality of evidence.	Metrics, corrections, access rights and retention periods.

An evidence pack may itself contain incorrect or incomplete information. Check source reliability, changes and access rights. Distinguish access to confidential information from review needs. Set retention periods according to law, policy and purpose.

Define each sign-off, such as preparation, review or approval within a stated scope. Do not infer that everyone has accepted every risk. Evidence supports assessment by management, those charged with governance and auditors, but does not guarantee reliance on controls. [8]

9 Stop contain and restart

Define stop triggers, alert recipients, stopping authority and backup channels before deployment. Examples include incomplete data, actions beyond permissions, correlated errors, cumulative volumes outside boundaries or failed alerting. Criteria must fit the entity's risks; there is no single numerical limit suitable for every system.

Table 9 Stop contain and restart

Stage	Actions	Reasons and records
Detect and suspend	Suspend affected activities; restrict permissions or volume.	Time, version and the trigger for stopping.
Determine affected scope	Identify transactions and periods from the point at which errors may have begun.	Known facts, uncertainties and the lookback method.
Contain and correct	Use fallback arrangements; correct data or transactions under proper authority.	Containment, corrections and responsible owners.
Communicate	Notify relevant parties according to their duties and applicable requirements.	Recipients, time, content and reasons for the notification scope.
Retest	Test the original cause, related cases and stopping capability.	Remediation results, new version and remaining issues.
Approve restart	Set scope, conditions and enhanced monitoring.	Authority, reasons, review date and limitations.

Restart does not follow automatically from a provider's announcement that an issue has been fixed. The process owner must assess the effects on the data and controls actually in use. The authorized person decides under the DOA, taking account of test results and remaining risk.

Review major incidents and near misses. Update knowledge, manuals, rules, role allocation and review capacity. Correcting individual outputs alone may leave the cause of recurring errors unresolved.

10 Hypothetical expected credit loss estimate

This case illustrates decision points for an entity and transactions within the scope of TFRS 9. It is not an expected credit loss (ECL) calculation formula or a conclusion for every SME. Those responsible must first check the applicable accounting framework and reporting period. [7]

AI may assist with segmentation, scenario analysis and draft explanations. The decision maker considers methods, assumptions, forecasts and indicators of bias. TFRS 9 paragraphs 5.5.17–5.5.18 concern unbiased, probability-weighted measurement, the time value of money, and reasonable and supportable information about past events, current conditions and forecasts. A longer historical dataset is not always better. [7]

Table 10 Hypothetical expected credit loss estimate

Decision point	Proposed controls	Reasons and evidence
Data and scope	Check receivables, period, completeness and segmentation.	Source records and the effects of missing data.
Methods and assumptions	Assess suitability for the portfolio and current conditions.	Supporting information and method limitations.
Material scenarios	Consider alternatives and sensitivity; avoid anchoring on the first answer.	Reasons for weights and information contradicting the proposal.
Review the estimate	A person with the information and competence challenges it before use.	Checks, corrections and remaining concerns.
Approve and post	Apply the DOA and actual permissions for each activity.	Approval scope, calculation version and accounting entries.

Use HITL at important decision points where people can actually review, and HIC for policy, boundaries and method changes. Choosing a loop does not replace applicable measurement and disclosure requirements.

TSA 540 (Revised) is an auditing standard for accounting estimates, including consideration of management bias. It does not require an entity to retain every AI answer. Evidence in this example is proposed to support review, not to guarantee auditor reliance. [8]

11 Hypothetical high volume receipt matching

The aim is to match receipts to receivables correctly and reduce rework, without inadvertently expanding permission from proposing a match to clearing balances or making payments. Distinguish control totals, payer identity, references and periods. Matching aggregate totals can still conceal allocation to the wrong account.

Table 11 Hypothetical high volume receipt matching

Process step	Oversight and controls	Evidence
Import transactions	Check sources, period, duplicates and completeness.	Control totals and dataset identifiers.
Propose or match within scope	Use approved rules and limits on the aggregate exposure.	Rule version, proposed matches and reasons.
Review exceptions	Human review before activities requiring a decision.	Unclear items, evidence and decision maker.
Execute permitted actions	Check permissions to post or close items.	Actual transactions and the authority used.
Monitor results	Check trends and sample items the system has passed.	False passes, correlated errors and backlog.
Stop and look back	Suspend when criteria are wrong or monitoring is unavailable.	Affected scope, corrections and retest results.

> HOTL requires evidence that alerts reach the responsible person and intervention occurs before unacceptable consequences. Monthly checking alone is insufficient. Specific items handled through HOOTL still need tested boundaries and an owner responsible for monitoring.

Metrics Track incorrect matches that pass undetected, duplicates, wrong periods, backlog, alert-to-stop time and the extent of harm before detection. Read these alongside independent checking: system confidence or a high match rate is not, by itself, evidence of accuracy.

The RCM must connect with risks to transactions and reporting, such as accuracy, completeness or cutoff, as relevant. Avoid applying the term occurrence to balances without considering what is being tested. A bank reconciliation alone does not cover every financial statement assertion.

12 Measure value and expand on evidence

Do not assume that knowing how to use AI creates a lasting advantage, or that fees for responsibility will always be high. DHISORA's position should center on connecting finance, risk, authority and processes to improve decisions and execution, with AI as a supporting tool.

Table 12 Measure value and expand on evidence

Dimension	Baseline before use	What to monitor after use
Efficiency	Preparation and review time; total cost.	Net time, including controls and corrections.
Quality	Errors, rework and unresolved issues.	Independent checks and errors that pass undetected.
Decision making	Previously retrievable information and reasons.	Completeness of information, challenges and rationale.
Risk	Previous affected scope and detection capability.	Alert-to-stop time and residual risk.
People's readiness	Competence, workload and authority.	Do reviewers have enough time and effective permissions?
Change	Previous versions, permissions and owners.	Change effects and the need for retesting.

A low override rate does not prove good results or an absence of human judgment. Consider independent checks, workload and evidence together. Reporting should show improvements, uncertainties and unresolved issues.

> Expanding use Examine one process → design and pilot within boundaries → assess results and decide on deployment → monitor before expanding. Timing depends on data, readiness and risk; results should not be guaranteed within a fixed number of days.

Service scope Offer reasoned assessment and design, with implementation within the agreed scope. References to COSO, ISO or AI-native status alone do not establish value. Responsibilities and professional licenses must match the work actually accepted.

Template 1 Use Case Card

> Purpose Define one AI activity clearly before assessing risk. Name the responsible people and attach supporting references. Illustrative fields are not the entity's approval criteria.

Field	Information to complete
Identifier and status	Use case ID / version / date / preparer [Complete]
Objective and baseline	Problem, desired outcome, metrics and current costs [Complete]
Process and scope	Entity, function, period, transaction types, volume and users [Complete]
Permitted and prohibited activities	Distinguish proposing, matching, posting, clearing balances, payments and reserved matters [Complete]
Data and systems	Sources, confidentiality classification, provider, version and permissions [Complete / reference records]
Owner and oversight	Process owner, reviewer, decision maker, approach and rationale [Complete]
Requirements and linked records	Applicable accounting framework or law / Risk IDs / Control IDs [Complete / identify who checked applicability]
Disposition	Develop further, restrict scope, suspend, or decline use / reasons / authority [Complete]

Quality check Can a reader tell what AI may actually do, where data comes from and who is responsible? Do not use AI automation as a substitute for specifying activities and permissions.

The inventory should also include actual use that has not been approved, so it can be assessed and addressed. Inclusion in the inventory does not constitute approval.

Template 2 Risk and Control Matrix

> Purpose Connect objectives with risks, controls, owners, testing and evidence. Use this template for one risk, then link controls that may address several risks.

Field	Information to complete
Identifiers and objective	Use case ID / Risk ID / relevant objective or assertion [Complete]
Cause event and impact	Data, system, people or provider / potential failure / business effects [Complete]
Inherent risk assessment	Likelihood, impact, speed, affected scope and supporting information [Complete / identify unknowns]
Control	Control ID / process point / prevention, detection, containment or correction [Complete]
Method and operator	Who does what, when, using which information, criteria and permissions? [Complete]
Testing and evidence	Tester, method, frequency, scope, expected results and storage location [Complete]
Residual risk assessment	Actual results, open issues, residual risk and rationale [Complete]
Decision and follow-up	Remediation owner, due date, risk acceptance authority where needed, and review date [Complete]

Hypothetical example R-MATCH-01: allocation to the wrong payer despite matching totals → C-ID-01: check payer and references → C-SAMPLE-01: independently sample items the system has passed → monitor correlated errors and suspension triggers.

Do not lower risk merely because HITL or an approver has been named. Control design and operating results must support the assessment.

Template 3 Authority and Oversight Record

> Purpose Extend the DOA to cover AI use without implicitly changing existing authority. Use one record per activity or decision, linked to policy, signing authority and actual permissions.

Field	Information to complete
Activity and references	Use case ID / decision type / DOA clause / version [Complete]
Authority and scope	Role / permitted decisions / financial limits or conditions / reserved matters [Complete]
Preparation review and execution	Roles / competence / cover for absence [Complete]
System and user permissions	Read, propose, change rules, post or perform other actions / prohibitions [Complete / reference access record]
Oversight and timing	HITL, HOTL or HOOTL in which cycle / who exercises HIC / response time [Complete]
Segregation and safeguards	Conflicting roles / supplementary review / responsible owner [Complete]
Exceptions stopping and restart	Escalation criteria / who may suspend / restart authority / channels [Complete]
Approval and review	Scope approver / date / condition expiry / changes requiring review [Complete]

Quality check Do actual permissions match authority? Can the alert recipient stop the process or only escalate? After a rule is approved, which transactions still require separate approval under applicable requirements?

Consider the substance and impact of exceptions. Routine data correction is not necessarily risk acceptance, and different sign-offs do not all mean the same thing.

Template 4 Judgment Record

> Purpose Record reasons for significant decisions according to risk. Use for estimates, exceptions, scope changes or matters where AI proposes a result and a person must choose. This is not an instruction to record every answer.

Field	Information to complete
Decision and scope	Decision ID / Use case ID / period / matter to decide [Complete]
Information and unknowns	Sources, versions, evidence, reliability and missing information [Complete / reference records]
AI proposal and limitations	Material output, known methods or assumptions, and matters not yet verifiable [Complete]
Material alternatives and challenges	Options affecting the conclusion / information contradicting the proposal [Complete]
Checks and changes	What was checked against which source / what was changed or rejected / why [Complete]
Conclusion and reasons	What was chosen, why, and subject to which limitations and conditions? [Complete]
Responsible people and sign-offs	Preparer / reviewer / approver under which DOA clause / date and time [Complete separately by role]
Follow-up and retention	Open issues, owner, due date, retention period and access rights [Complete]

> Insufficient wording “Reviewed and agreed with AI” does not explain enough. Identify the information checked, material issues and reasons for the choice so that others can understand later.

If residual risk is being accepted, identify its scope, risk, rationale and authority separately. Do not infer acceptance from a preparer's or reviewer's signature.

Template 5 Test Plan and Results

> Purpose Set criteria in advance and retain reproducible test results. Use one record per test case and summarize coverage in the overall plan. Testers should have competence and independence appropriate to the substance of the work.

Field	Information to complete
Identifiers and scope	Test ID / Use case / Control ID / data and system versions / test period [Complete]
Objective and cases	What is being tested / normal, abnormal, boundary or failure cases [Complete]
Criteria and verification basis	Predefined criteria / expected results / independent verification source [Complete / date criteria were approved]
Method and coverage	Steps, sample size, selection method and rationale / limitations [Complete]
Environment and safety	Separation from live accounting / permissions / approval for simulated errors [Complete]
Actual results and evidence	Results against criteria / storage / tester / date and time [Complete]
Deficiencies and remediation	Substance, impact, owner, due date and affected conditions [Complete]
Retest and conclusion	Retest results / remaining issues / reviewer / supported conclusion scope [Complete]

The overall plan should cover data, logic, permissions, changes, alerts, stopping, fallback and human workload according to risk. Inspecting design documents alone is not equivalent to testing operation.

Changes to criteria require reasons, authority and retesting. Do not make earlier results appear to pass by deleting limitations or silently adjusting criteria.

Template 6 Go Live Decision Record

> Purpose Enable an authorized person to decide on deployment from evidence actually considered, rather than simply confirming that documents are complete. This assessment does not replace system certification or auditor assurance.

Field	Information to complete
Proposed deployment scope	Use case ID / version / function, period, activities and volume [Complete]
Requirements and authority	Requirements checked / checker / authorized person under the DOA [Complete]
Test results and limitations	Test summary / criteria met or unmet / evidence [Complete / reference records]
Operational readiness	People, permissions, manuals, evidence, alerts, stopping and fallback [Complete / identify gaps]
Residual risk	Risk, evidence, acceptance rationale and authority [Complete / identify mandatory matters that cannot be waived]
Decision	Approve / approve with conditions / defer / decline use, with reasons [Select and explain]
Conditions and validity	Limits, boundaries, oversight owner, expiry and suspension triggers [Complete where applicable]
Follow-up and sign-off	Review date, metrics, report recipients / decision maker, date and time [Complete]

> Final question before deployment What happens if data is wrong, alerts disappear or the reviewer is absent? Who can stop the process, and which activities must revert to fallback arrangements?

Conditional approval is not a route around mandatory requirements or deficiencies that leave significant outcomes uncontrolled.

Template 7 Incident and Restart Record

> Purpose Connect the incident, containment, correction and restart in a reviewable record. Clearly state unknowns and revise the affected scope as further evidence emerges.

Field	Information to complete
Event and detection	Incident ID / Use case / date and time / alert and reporter [Complete]
Versions and affected scope	Data, model, rules and permissions / transactions, periods or affected parties [Complete / identify uncertainty]
Suspension and containment	What stopped, by whom and when / fallback and restricted permissions [Complete]
Impact and communication	Actual and potential effects / required recipients / date and time [Complete]
Cause and correction	Data, system, person or process / remedy, owner and due date [Complete]
Post-remediation tests	Test IDs / original and related cases / stopping / open issues [Complete]
Restart proposal and decision	Scope, version and conditions / authority and rationale / date and time [Complete]
Monitoring and closure	Enhanced monitoring, review date, actual results, lessons and closure owner [Complete]

Do not close an incident solely because an individual item was corrected. Check whether a common cause remains, earlier items are affected, and alerts and stopping actually work.

The incident owner and restart approver must have authority appropriate to the substance and risk. They need not always be the person who approved the original transaction.

Template 8 Management Monitoring Review

> Purpose Summarize actual results against the baseline, open issues and decisions needed at review. Set frequency according to risk and change. Periodic reporting does not replace alerts that must be timely.

Field	Information to complete
Scope and reporting period	Use cases / process / period / report owner [Complete]
Value against baseline	Net time, total cost, quality and rework / measurement method [Complete]
Quality and independent checks	Errors passing undetected, sample results and affected scope [Complete / reference evidence]
Oversight and workload	Do reviewers have information, time and authority? / backlog / responses [Complete]
Incidents and exceptions	Significant incidents, alert-to-stop time, open issues and remediation dates [Complete]
Changes and requirements	Versions, data, permissions and providers / matters requiring testing or review [Complete]
Residual risk and unknowns	Latest assessment, information limitations and escalation matters [Complete]
Decision and follow-up	Maintain, restrict, expand, suspend or correct / reasons, owner and review date [Complete]

Reporting to those charged with governance should identify matters material to objectives and risks. Do not report only prompt counts or the percentage of work performed by AI. Specify decisions needed, the relevant authority and available options.

Use actual results to update use cases, the RCM, DOA, manuals and test plans. Record benefits alongside the additional costs of control.

Governance and control references

Official sources and status checked as of 10 October 2026. References inform the concepts; they do not represent a declaration of compliance with every requirement.

1 COSO GenAI 2026

Guidance on applying internal control to GenAI, not a law or certificate. Used to support design.

Open official source

2 NIST AI RMF and GenAI Profile

The AI Risk Management Framework is voluntary. NIST's page indicates a revision is in progress. Use published versions and check status before describing them as requirements.

Open official source

3 ISO IEC 42001 2023

AI management system requirements. Referencing the standard does not establish that an organization is certified.

Open official source

4 EU AI Act

A law whose scope, roles, system categories and application dates must be checked. Not every finance activity is high-risk.

Open official source

5 Bank of Thailand AI policy 12 September 2025

An issued policy for financial institutions, specialized financial institutions and payment service providers within its scope, not a requirement for every company.

Open official source

6 IIA Statements of Position

Current sources on the Three Lines Model and internal audit's role. The IIA page identifies the revised publication date as 8 July 2026.

Open official source

Professional and human oversight references

Official sources and status checked as of 10 October 2026. References inform the concepts; they do not represent a declaration of compliance with every requirement.

7 TFAC Financial Reporting Standards

Check the framework applicable to the entity and the effective reporting period. The ECL example refers to TFRS 9 where applicable.

Open official source

8 TFAC Auditing Standards

Check applicable standards, including TSA 330 and TSA 540 (Revised). They govern auditors' work; they are not AI system certificates.

Open official source

9 European AI HLEG 2019

The Ethics Guidelines for Trustworthy AI discuss HITL, HOTL and HIC as ethical guidance, distinct from the AI Act's legal provisions.

Open official source

10 Australian DTA Key Terms

Definitions of HITL, HOTL and HOOTL in the Australian public sector context. Used to inform design terminology, not extended into universal rules.

Open official source

11 Federal Reserve SR 26 2

The 2026 guidance replaces SR 11-7 and addresses model risk in banking. Check its scope; do not present it as a general requirement for GenAI or agentic AI.

Open official source

ISO 31000 and ISO/IEC 23894 provide risk management guidance, whereas ISO/IEC 42001 specifies management system requirements. Do not claim to have passed ISO without defining the scope and providing certification evidence.

Before using a law or specific guidance in a decision, recheck the latest version, organizational roles, systems within scope and applicable dates.

Editorial review and content direction

This edition preserves the judgment, human oversight, DOA, risk and process themes of Claude's original material and reorganizes them into an operating guide. It does not edit the original file or provide a sentence-by-sentence redline. Added material covers readiness criteria, ongoing operation, stopping and restart, and templates that connect evidence.

Table 13 Editorial review and content direction

Issue reviewed	Approach in this edition
Evidence and reliance	Replace guarantees of results with information supporting assessment within each role.
Authority and automated rules	Rule approval does not replace transaction authority reserved by applicable requirements.
Loops and reviewer quality	No fixed ladder: HIC spans other approaches, and HITL must enable effective review.
Governance roles	Distinguish management, second-line functions, internal audit and the audit committee.
Risk and evidence	Sign-off alone does not justify a lower rating; not every AI answer needs retention.
Reference versions and scope	Use the issued BOT policy and current IIA publications; distinguish SR 11-7 from its replacement.
Implementation	Add eight templates and criteria for testing, go-live, incidents and restart.

Role of the website The executive synthesis is the starting point. Four articles explain the core concepts; the playbook article connects them to practice. This file can be adapted into entity-specific documentation. Reduce theoretical repetition and avoid positioning DHISORA as an AI certifier or an expert in every discipline.

DHISORA's direction Focus on accountable decisions and controls in finance. Advantage must be demonstrated through contextual understanding, connections across disciplines and actual operating results, not through artificial scarcity, difficult access or marketing labels alone.