Gamut Auditor & Assurance Manual

Customer-ready launch manual

Version	Launch version
Audience	External auditors, internal audit, assurance reviewers, compliance reviewers, risk committees, and independent assessors

Table of Contents

About This Manual

This manual explains how auditors and assurance reviewers use Gamut to review AI governance records. It focuses on read-only review, evidence inspection, findings, testing, workpaper packs, audit trail review, board reporting, and certification readiness.

Important note: Gamut helps organise and present assurance evidence. It does not replace auditor judgment, legal advice, regulatory advice, or formal certification processes.

Auditor Access Model

Auditor access is designed for independent review.

An auditor can typically review:

- Dashboard summaries.
- AI system records.
- Framework assessments.
- Risk Register.
- Findings Register.
- Evidence Tracker.
- Testing Centre.
- Workpaper Packs.
- Reports.
- Policies.
- Model Cards.
- Audit Trail.

Auditors are normally read-only. They should not create, edit, delete, approve, export, invite users, or spend AI analysis quota unless the organisation gives them a different role.

If you need to request changes, ask the accountable owner or administrator to update the record.

Invitation And First Login

An administrator may invite you to Gamut as an external auditor.

The usual first-login flow is:

1. Open the invitation link.
2. Create or confirm your account details.
3. Set or change your password.
4. Verify your email if prompted.
5. Enroll multi-factor authentication.
6. Open the assigned workspace.

Auditor access may be time-limited. If your link or access has expired, contact the organisation that invited you.

Multi-Factor Authentication

External auditor sessions require MFA before accessing data.

During MFA setup:

- Add the Gamut account to your authenticator app.
- Enter the current 6-digit code.
- Save backup codes securely.

If you lose access to your authenticator and backup codes, contact the organisation's Gamut administrator.

Read-Only Review Principles

As an auditor or assurance reviewer:

- Treat the workspace as the review record.
- Do not rely only on dashboards.
- Follow conclusions back to system records, intake, scores, evidence, tests, findings, and remediation.
- Look for missing rationale, weak evidence, stale records, and unclear ownership.
- Record questions outside Gamut unless your organisation has given you a permitted comment or review workflow.
- Ask accountable owners to make corrections rather than attempting to edit records.

Workspace Orientation

Start by confirming the workspace scope.

Check:

- Organisation or client name.
- Assessment title.
- Assessment period or review timing.
- Systems in scope.
- Frameworks in scope.
- Report or workpaper pack purpose.
- Whether the workspace is complete, draft, or still being updated.

If the scope is unclear, request clarification before relying on any conclusions.

Dashboard Review

Use the Dashboard to understand the high-level status of the workspace.

Review:

- Overall posture.
- Progress indicators.
- High-risk areas.

- Open findings.
- Evidence gaps.
- Assessment completeness.
- Readiness signals.

The Dashboard is an entry point, not an audit conclusion. Use it to decide where to inspect next.

AI System Records Review

AI System Records shows the systems in scope.

For each system, check:

- The system has a clear name and identifier.
- The purpose is understandable.
- A business owner is named.
- Lifecycle stage is current.
- Vendor or deployment type is recorded.
- Data categories and deployment environment are clear.
- Risk tier is not missing or unexplained.
- The record aligns with intake and assessment scope.

Questions to ask:

- Are all material AI systems registered?
- Are retired, paused, pilot, and production systems clearly distinguished?
- Are high-impact systems reviewed with enough depth?

AI Use Case Intake & Approval Review

AI Use Case Intake & Approval records explain what each AI system is intended to do.

Review:

- Use case summary.
- System description.
- Users and affected persons.
- Data sources.
- Data classification.
- Geographies or jurisdictions.
- Regulatory exposure.
- Approval status.
- Reviewer notes.

Good intake should support risk routing and assessment scoping. If intake is thin, downstream scores may be unreliable.

Risk Tiering And Assurance Routing Review

Risk Tiering Engine and Assessment Plan & Assurance Routing help determine governance depth.

Review:

- Whether risk factors are populated.
- Whether autonomy, access, dependency, and harm exposure are credible.
- Whether risk tiering aligns with the use case.
- Whether high-risk routes led to deeper evidence, testing, findings, or approval requirements.
- Whether manual overrides are explained.

If risk routing appears too low for the use case, request rationale and supporting evidence.

Framework Assessment Review

Gamut may contain assessments for frameworks such as GTSAF, ACRS, ATF, MAESTRO, EU AI Act, NIST AI RMF, ISO/IEC 42001, ISO/IEC 42005, and NAGF.

When reviewing any framework:

- Confirm the framework is in scope.
- Check whether the assessment is complete enough for the intended use.
- Review low-scoring or non-compliant areas.
- Review not-applicable rationale.
- Check notes for scoring rationale.
- Confirm evidence exists for important claims.
- Confirm findings or remediation exist for material gaps.

Do not treat a score as sufficient without evidence and rationale.

GTSAF Assurance Review

GTSAF is often the central control assurance view.

Review:

- Control conformance.
- Evidence quality.
- Ownership clarity.
- Assurance and audit-readiness signals.
- Shared-responsibility notes.
- Control gaps.

Important questions:

- Are controls scored consistently?
- Is evidence specific to the control?
- Are weak controls connected to risks, findings, or remediation?

- Are not-applicable answers justified?
- Is there enough evidence for external reliance?

EU AI Act Review

For EU AI Act review, inspect:

- Applicable role and risk route.
- Classification rationale.
- High-risk or GPAI-related obligations where relevant.
- Technical documentation.
- Human oversight.
- Transparency.
- Risk management.
- Post-market monitoring.
- Non-applicability rationale.

Regulatory interpretation should be confirmed by qualified legal or compliance professionals.

NIST AI RMF Review

For NIST AI RMF review, inspect:

- Govern, Map, Measure, and Manage coverage.
- Whether risk management activities are operational, not merely documented.
- Whether evidence supports the stated maturity.
- Whether high-risk systems have appropriate controls and review cadence.

ISO Readiness Review

Where ISO-oriented modules are enabled, use them to assess readiness.

Review:

- Scope and context.
- Leadership and governance.
- Planning and support.
- Operational controls.
- Performance evaluation.
- Improvement actions.
- Evidence mapped to clauses or requirements.

Gamut readiness views are planning aids. Formal certification requires the appropriate certification process.

NAGF Review

For NAGF review, inspect:

- Governance records.
- Data protection considerations.
- Sector-specific exposure.
- Evidence supporting regulatory readiness.
- Gaps that require legal, compliance, or sector review.

Confirm legal obligations with qualified legal or regulatory advisors.

MAESTRO Review

Use MAESTRO to review threat and risk analysis.

Check:

- Threat coverage.
- Likelihood and impact scoring.
- Cross-layer risks.
- Mitigation quality.
- Links to findings, risks, or remediation.

Threat analysis should be proportionate to system criticality and exposure.

ATF And Agentic Assurance Review

If AI agents are in scope, review the agentic governance records.

Check:

- Agent registration.
- Named human and security owners.
- Identity profiles.
- Access matrix.
- Tool permissions.
- Data flows.
- Approval gates.
- Incident playbooks.
- Red-team tests.
- Gateway policy decisions.
- ATF maturity level and promotion rationale.

Do not rely on an agent readiness claim unless ownership, access, approvals, testing, and evidence are complete.

Discovery Assurance Review

If Discovery is enabled, review it to understand whether the AI inventory is supported by observed signals or only by voluntary registration.

Discovery review should cover:

- Monitor.
- Priorities.
- Playbooks.
- Rules.
- Security.
- Coverage.
- Assets.
- Review queue.
- Workflow gaps.
- Sources.
- Discovery Assurance Report where generated.

Monitor Review

Use Monitor to inspect recent discovery activity, alerts, drift, source attention, new signals, and new artifacts.

Review:

- Whether collectors were run recently enough for the review period.
- Whether high-risk, unsanctioned, or repeated signals remain open.
- Whether active alerts have been acknowledged, resolved, suppressed, or converted appropriately.
- Whether alert suppression includes a reasonable rationale.
- Whether alerts converted into risks, findings, or evidence requests can be traced to those records.

Unresolved high-risk alerts should be treated as open assurance questions.

Priorities Review

Use Priorities to confirm that the organisation is triaging the highest-value discovery items first.

Review:

- Urgent or high-severity lanes.
- Items related to unsanctioned tools, regulated data, external exposure, repeated activity, or missing owners.
- Whether priority items were reviewed and routed to intake, risks, findings, evidence, or remediation.
- Whether stale priority items remain unassigned.

If a high-risk discovery priority has no downstream action, request reviewer rationale.

Playbooks Review

Playbooks show guided responses generated from discovery signals and alerts.

Review:

- Which playbooks are waiting or were executed.
- The severity and rationale for each playbook.
- Whether execution created appropriate linked records.
- Whether created risks, evidence requests, or findings were completed after the playbook ran.

A playbook execution is not final assurance evidence by itself. Inspect the resulting records.

Rules Review

Rules determine how discovery signals are normalised and grouped.

Review:

- Whether custom rules are precise enough to avoid hiding distinct systems.
- Whether aliases, domains, or provider names map logically to the canonical tool or system.
- Whether broad rules could cause undercounting of AI systems.
- Whether rules explain unusual grouping in Assets or Review queue.

If rules are overly broad, ask the organisation to demonstrate that separate tools or use cases were not incorrectly merged.

Security Review

Security shows whether discovery is operating under controlled collection expectations.

Review:

- Guardrail status.
- Whether discovery fails closed when source health is poor.
- Whether source records appear to avoid storing secrets.
- Which sources are live, manual, attention-needed, or failing.
- Whether connector status aligns with least-privilege and read-only expectations.

If sensitive data, credentials, raw prompts, raw completions, or unnecessary source content appear in discovery records, escalate the issue through the organisation's confidentiality process.

Coverage Review

Coverage shows whether the discovery programme is likely to find shadow AI across expected surfaces.

Review:

- Overall coverage score or posture.
- Covered source categories.
- Missing source categories.
- Failing sources.
- Manual-only areas.
- Coverage notes describing blind spots.

Low coverage means the AI inventory may be incomplete. Do not accept "no shadow AI found" as a strong conclusion when major source categories are missing or failing.

Assets Review

Assets group discovery artifacts into suggested AI assets.

Review:

- Whether suggested assets have been reconciled.
- Whether promoted assets became AI System Records and intake items.
- Whether duplicate or related artifacts are handled consistently.
- Whether artifacts with credible AI usage remain unowned or unrouted.

Trace a sample of promoted assets from artifact to system record, intake, risk route, and evidence.

Review Queue Review

The Review queue contains discovered candidates that require triage.

For sampled candidates, check:

- Source and source type.
- Confidence and risk indicators.
- Owner, department, user, domain, provider, or application context.
- Evidence excerpt or metadata summary.
- Status and reviewer decision.
- Dismissal rationale where applicable.
- Links to system, intake, risk, evidence, finding, or remediation records.

Dismissed candidates should not be accepted without a rationale, especially when they involve high-risk tools, sensitive data, external users, or repeated signals.

Workflow Gaps Review

Workflow gaps show discovery items that have not been routed through governance.

Review gaps such as:

- Candidate with no system record.
- Asset with no intake.
- Missing owner or department.
- High-risk signal with no evidence request, risk, finding, or remediation.
- Failing or stale source.
- Missing coverage area.
- Review queue item overdue.

Workflow gaps are often the clearest evidence that the governance process is incomplete.

Sources Review

Sources define where discovery signals come from.

Review each material source for:

- Source type and connector.
- Collection mode.
- Owner and owner email.
- Scope, tenant, organisation, or workspace hint.
- Cadence and scheduled-pull status.
- Monitored apps or domains.
- Coverage note.
- Health status.
- Last run or recent activity where shown.
- Manual import or sample input handling where applicable.

Sources should be scoped, owned, current, and documented. A source with no owner, unclear scope, failing status, or no coverage note should be treated as an assurance limitation.

Discovery Assurance Report Review

If a Discovery Assurance Report is provided, confirm:

1. The report generation date.
2. The workspace and scope.
3. Sources and coverage.
4. Open candidates and review queue status.
5. Artifacts and promoted assets.
6. Priorities and alerts.
7. Workflow gaps and limitations.
8. Whether high-risk discovery items were routed into the formal governance workflow.

The report is useful evidence of discovery operations, but the auditor should still sample the live workspace records where access permits.

Risk Register Review

The Risk Register should show the major AI risks and their treatment.

Review:

- Risk title and description.
- Category or framework link.
- Inherent and residual risk.
- Owner.
- Treatment plan.
- Review date.
- Status.

- Acceptance or escalation rationale.

Questions to ask:

- Are high-risk systems reflected in the risk register?
- Are treatment plans specific and time-bound?
- Are accepted risks formally justified?
- Are overdue risks visible?

Evidence Tracker Review

The Evidence Tracker is central to assurance review.

For each important evidence record, check:

- What claim or control it supports.
- Whether it is linked to the correct system, framework, control, finding, or risk.
- Owner.
- Due date.
- Status.
- Quality rating.
- Review notes.
- Whether evidence is current.
- Whether evidence is specific enough for the claim.

Evidence weaknesses include:

- Generic policy evidence for a system-specific claim.
- Old evidence with no review date.
- Evidence without an owner.
- Evidence not linked to the relevant control.
- Screenshots or documents without context.
- Evidence that describes intended design but not operating effectiveness.

Testing Centre Review

The Testing Centre shows active control verification.

Review:

- Test objective.
- Control or requirement being tested.
- Test method.
- Sample or scope.
- Test result.
- Exceptions.

- Owner.
- Test date.
- Retest requirement.
- Links to findings or remediation.

Testing is important where documentation alone is not enough. Failed or inconclusive tests should be linked to findings or remediation actions.

Findings Register Review

The Findings Register should show material issues and their resolution status.

Review:

- Finding title and description.
- Severity.
- Status.
- Owner.
- Due date.
- Evidence summary.
- Root cause.
- Management response.
- Closure evidence.
- Validation notes.

Important questions:

- Are high-severity issues clearly described?
- Are findings assigned to accountable owners?
- Are due dates realistic?
- Are closed findings supported by closure evidence?
- Are repeated findings linked to broader remediation?

Remediation Roadmap Review

The Remediation Roadmap should turn risks, findings, test failures, and assessment gaps into managed actions.

Review:

- Action title.
- Priority.
- Owner.
- Due date.
- Status.
- Dependencies.

- Closure criteria.
- Evidence required for closure.

Weak remediation actions are vague, unowned, overdue, or impossible to validate.

Model Cards Review

Model Cards support transparency and system understanding.

Review:

- Intended use.
- Limitations.
- Evaluation information.
- Data considerations.
- Bias and fairness notes.
- Monitoring arrangements.
- Ownership.
- Review status.

Model Cards should align with system records and intake.

Policy Review

Where policy documents are available, review whether they support the control environment.

Check:

- AI governance policy.
- Risk management standard.
- Data governance procedure.
- Impact assessment procedure.
- Incident management procedure.
- Model lifecycle procedure.
- Third-party management procedure.
- Human oversight guidance.
- Bias testing procedure.

Policies should be approved, current, and connected to actual operating records.

Audit Trail Review

The Audit Trail shows activity and changes in the workspace.

Use it to review:

- Score changes.
- Evidence uploads or deletions.

- Evidence request changes.
- Control test changes.
- Finding changes.
- Intake changes.
- Workpaper pack changes.
- Agentic governance changes where in scope.

Audit trail review helps validate timing, accountability, and change history. Use filters and search where available.

Workpaper Packs

Workpaper Packs package governance records for review.

A pack may include:

- Assessment summary.
- System records.
- Intake records.
- Risk routing.
- Framework scores.
- Evidence.
- Tests.
- Findings.
- Remediation.
- Reports and limitations.

When reviewing a pack:

- Confirm the pack scope.
- Confirm the generation date.
- Check whether it is workspace-wide or system-specific.
- Review assumptions and limitations.
- Trace important claims back to evidence.
- Request updates if the live workspace has changed since pack generation.

Board Dashboard Review

The Board Dashboard is designed for executive oversight.

Review:

- Key posture indicators.
- High-risk systems.
- Critical or high findings.

- Evidence coverage.
- Overdue actions.
- Decisions needed.
- Board attention items.

The Board Dashboard is useful for governance committees, but it is not a substitute for detailed testing and evidence review.

Reports And Exported Outputs

Reports reflect the records available at the time they are generated.

Export options may include Trustworthy AI PDF and Trustworthy AI Word, depending on role and plan.

Before relying on a report:

- Confirm report scope.
- Confirm report date.
- Confirm included frameworks.
- Check whether key evidence and findings are current.
- Review assumptions and limitations.
- Confirm draft or final status.
- Confirm the organisation has approved any external sharing of confidential evidence, personal data, or draft conclusions.

If the underlying workspace changes, regenerate or update the report.

Assessment History

Assessment History helps reviewers understand prior states.

Use it to:

- Review earlier assessments.
- Compare progress.
- Identify recurring gaps.
- Validate whether remediation improved posture.

History is especially useful for recurring audits and periodic governance reviews.

Assessment Baselines & Delta Tracking Review

Assessment Baselines & Delta Tracking compares current posture against a prior baseline.

Review:

- Improved areas.
- Declined areas.
- Unchanged high-risk gaps.
- Evidence supporting claimed improvement.

- Whether baseline and current assessments are comparable.

A claimed improvement should be backed by changed evidence, tests, findings, or remediation status.

Automated Gap Narratives

Automated Gap Narratives describe control gaps and expected evidence in plain language.

Use them as drafting aids, not final conclusions.

Before relying on a narrative:

- Check that the underlying score is correct.
- Check that evidence supports the wording.
- Confirm whether the baseline is relevant.
- Confirm that unchanged areas were actually reassessed.

Cross-Framework View

The Cross-Framework View shows how weaknesses affect multiple frameworks.

Use it to:

- Identify common control dependencies.
- Avoid duplicate remediation.
- Understand where one gap creates several compliance issues.
- Prioritise actions with the broadest assurance benefit.

Cert Readiness

Cert Readiness provides a high-level readiness signal.

Review:

- Overall readiness score.
- Framework-specific gaps.
- Evidence readiness.
- Open findings.
- Remediation status.
- Control maturity.

Readiness is not certification. Treat it as a planning indicator and validate with formal assessment criteria.

Global Search

Use Global Search to locate records across the workspace.

Search can help find:

- Systems.
- Intake records.

- Risks.
- Evidence.
- Tests.
- Findings.
- Remediation.
- Policies.
- Model Cards.
- Reports.

Search is useful when tracing a claim or responding to a review question.

Assurance Review Method

A practical review sequence:

1. Confirm workspace scope.
2. Review system inventory.
3. Review intake and risk routing.
4. Review Discovery coverage, sources, queue, priorities, assets, and workflow gaps where enabled.
5. Inspect framework scores and notes.
6. Review risks and findings.
7. Inspect evidence for high-risk and low-scoring areas.
8. Review control tests.
9. Review remediation status.
10. Inspect audit trail for major changes.
11. Review workpaper packs and reports.
12. Summarise open questions and evidence gaps.

What To Request From The Organisation

If records are incomplete, request:

- Missing system owner.
- Missing intake rationale.
- Risk route explanation.
- Control scoring rationale.
- Not-applicable rationale.
- Current evidence.
- Evidence owner and review date.
- Test procedure or sample.
- Closure evidence for findings.

- Remediation owner and due date.
- Approval evidence for high-risk agent actions.
- Updated workpaper pack after changes.

Common Issues

I cannot edit a record:

- Auditor access is normally read-only. Ask the organisation to update the record.

I cannot export:

- Auditor access may not include export rights. Ask the organisation to generate the required pack or report.

I cannot see a workspace:

- Your invitation may not include that workspace, or access may have expired.

The dashboard looks clean but evidence is missing:

- Dashboards reflect available data. Inspect evidence and testing before relying on the posture.

A control is marked not applicable without rationale:

- Request rationale, decision owner, and supporting basis.

A finding is closed without closure evidence:

- Request closure evidence and validation notes.

Auditor Checklist

Before concluding a review, confirm:

- Workspace scope is clear.
- Systems in scope are registered.
- Intake records are complete enough to support risk routing.
- High-risk systems are assessed with appropriate depth.
- Discovery coverage, sources, review queue, priorities, assets, and workflow gaps have been reviewed where enabled.
- Framework scores include rationale.
- Not-applicable answers include rationale.
- Evidence is current, specific, and linked.
- Tests support key operating-effectiveness claims.
- High and critical findings are owned and tracked.
- Closed findings have closure evidence.
- Remediation actions have owners, dates, and closure criteria.
- Model Cards are complete where relevant.
- Agentic systems have owner, access, tool, approval, testing, and evidence records.

- Audit trail does not show unexplained late changes.
- Workpaper packs and reports match the live workspace record.
- Open questions are documented for management response.

Comprehensive Feature Handbook

This handbook layer gives auditors and assurance reviewers a repeatable method for using Gamut without editing customer records.

Assurance Scoping

Confirm workspace, period, frameworks, systems, review purpose, and record maturity before relying on conclusions.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.
- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Inventory Completeness Review

Sample system records, discovery outputs, intake records, and retired or paused systems to assess completeness.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.
- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Evidence Sufficiency Review

Inspect evidence quality, relevance, recency, owner, linkage, and review notes.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.
- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Control Score Review

Trace high scores, low scores, not-applicable selections, and late changes back to evidence and rationale.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.
- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Risk And Finding Review

Check whether risks and findings are clear, owned, dated, treated, and validated at closure.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.

- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Discovery Assurance Review

Review sources, coverage, monitor, priorities, playbooks, rules, assets, queue, workflow gaps, and assurance report.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.
- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Agentic Assurance Review

Review agent register, identity, access, tools, data movement, approvals, red-team tests, Gateway decisions, and runtime readiness.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.
- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Report Reliance Review

Confirm generated reports and workpapers match the live workspace and do not overstate readiness.

Primary users: external auditors, internal audit, assurance reviewers, and independent assessors.

What this feature does:

- Captures the facts needed for accountable AI governance.
- Connects the record to later review, evidence, risk, reporting, or assurance work.
- Helps reviewers distinguish complete, partial, stale, and unsupported records.
- Gives owners a practical place to keep status, rationale, and next actions current.

Before you start:

- Confirm that the correct client and workspace are open.
- Confirm that edit mode is available if you need to make changes.
- Confirm that you have enough source information to complete the record without guessing.
- Prepare owner names, dates, scope, supporting evidence, and review rationale before saving material decisions.

Field guide:

- Scope: The workspace, system, framework, or report boundary being reviewed.
- Evidence: The proof inspected.
- Rationale: The reasoning supporting a score, decision, or conclusion.
- Exception: A gap, uncertainty, missing record, or inconsistency requiring management response.
- Reliance note: Whether the reviewer can rely on the record for the intended purpose.

How to use this feature:

1. Confirm scope and access limitations.
2. Use dashboards only to plan review, not to conclude.
3. Trace samples from summary to source records.
4. Inspect evidence and rationale for material claims.
5. Identify missing owners, stale dates, unsupported scores, and unlinked gaps.
6. Request corrections from accountable owners rather than editing records.
7. Document open questions and management responses.

Quality checks before you move on:

- Every conclusion can be traced to records.
- Evidence is current, specific, and linked.
- Material gaps have management action.
- Reports match the live workspace.
- Read-only independence is preserved.

Common problems and fixes:

- The record is too vague: Rename the item and rewrite the purpose so a reviewer can identify the real system, use case, action, or control.
- The owner is missing: Assign a named person or accountable function before relying on the record.
- Evidence is missing: Create an evidence request or attach a supporting record before increasing confidence or readiness.
- The status is stale: Update status, review date, and notes after material changes or review meetings.
- The record is not linked: Link the item to the relevant system, intake, risk, finding, test, evidence, or remediation record.

Example of a strong record:

A strong record names the system or activity clearly, identifies the owner, explains the business purpose, states the current status, links to evidence or related records, and records enough rationale for another reviewer to understand why the decision was made.

Example of a weak record:

A weak record uses a vague title, has no accountable owner, omits scope or dates, contains unsupported conclusions, or leaves the next action unclear. Weak records should be returned to the owner for completion before being used in reports.

Auditor Procedure Library

Use these procedures to perform read-only assurance review in a consistent way. The auditor's job is not to complete the customer's records, but to determine whether the records are complete, supported, current, scoped, and reliable enough for the intended conclusion.

Accept An Auditor Invitation

Scenario: An auditor has received access and must complete first login securely.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.

6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Confirm Review Scope

Scenario: The reviewer needs to understand what workspace, period, framework, systems, and reports are in scope.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?

- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Assess Workspace Maturity

Scenario: The auditor needs to decide whether the workspace is ready for reliance or still draft.

Records and views usually involved:

- Dashboard and workspace overview.

- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.

- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review System Inventory Completeness

Scenario: The auditor needs to determine whether material AI systems are registered.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.

2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.

- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Sample AI System Records

Scenario: A sample of systems must be checked for owner, purpose, lifecycle, vendor, data, and risk context.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?

- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Intake Quality

Scenario: Use case records need to support risk routing and assessment scope.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.

- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Risk Tiering

Scenario: The risk route must be proportionate and explainable.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.

- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review ACRS And Agentic Routing

Scenario: Agentic capability scores need to match autonomy, access, harm, and product context.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.

9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.

- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review GTSAF Scores

Scenario: Core control scores need rationale and evidence.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?

- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review EU AI Act Records

Scenario: EU role, classification, obligations, and non-applicability rationale must be checked.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.

- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.

- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review NIST AI RMF Records

Scenario: Govern, Map, Measure, and Manage activities must be operationally supported.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.

6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review ISO Readiness

Scenario: Readiness indicators must be treated as planning aids, not certification.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?

- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review NAGF Alignment

Scenario: Nigerian governance relevance and legal questions must be visible where applicable.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.

- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.

- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review MAESTRO Threat Analysis

Scenario: Threat scenarios and mitigations must be proportionate to exposure.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.

3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.

- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review ATF And Agentic Assurance

Scenario: Agents need ownership, identity, access, tools, approvals, testing, and runtime evidence.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?

- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Discovery Coverage

Scenario: Shadow AI discovery should have sources, coverage notes, review queue decisions, and workflow gap handling.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.

- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Discovery Monitor And Alerts

Scenario: Discovery alerts must be acknowledged, resolved, converted, or justified if suppressed.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.

- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Discovery Sources

Scenario: Sources must be owned, scoped, healthy, and documented with limitations.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.

9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.

- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Discovery Assets And Queue

Scenario: Candidates and assets must be confirmed, linked, promoted, dismissed with rationale, or routed.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?

- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Risk Register

Scenario: Risks need description, owner, treatment, residual view, status, and review date.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.

- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.

- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Evidence Tracker

Scenario: Evidence must be current, specific, linked, reviewed, and sufficient.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.

6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Testing Centre

Scenario: Tests must have objective, method, sample, result, exceptions, and links to findings.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?

- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Findings Register

Scenario: Findings must be clear, severe enough, owned, dated, responded to, and supported at closure.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.

- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.

- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Remediation Roadmap

Scenario: Actions must be specific, owned, time-bound, and testable.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.

3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.

- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Model Cards

Scenario: Model transparency records must align with system and intake records.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?

- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Policies

Scenario: Generated or uploaded policies must be current, approved, and connected to operation.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.

- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Audit Trail

Scenario: Major late changes, score changes, evidence changes, and exports must be explainable.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.

- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Workpaper Packs

Scenario: Packs must match live workspace scope and generation date.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.

- Auditor access remains read-only unless explicitly authorised otherwise.

Review Board Dashboard

Scenario: Executive summaries must not overstate readiness or hide evidence gaps.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Exported Reports

Scenario: Reports must match the source workspace and be approved for the audience.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.

- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Baseline And Delta Claims

Scenario: Improvements must be backed by changed evidence, tests, closure, or remediation.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.

8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.

- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Gap Narratives

Scenario: Narratives must reflect the underlying score and evidence.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?

- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Cross-Framework Mapping

Scenario: Shared controls must not be assumed to satisfy all frameworks without evidence.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.

- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.

- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Certification Readiness

Scenario: Readiness must not be represented as certification.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.

4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.

- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Global Search Results

Scenario: Search should be used to find duplicates, missing links, and hidden records.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?

- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Trace A Report Claim To Evidence

Scenario: A report statement must be followed back to system, intake, score, evidence, and action.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.

- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Investigate Unsupported High Scores

Scenario: High maturity or compliance scores lack adequate proof.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.

- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Investigate Not-Applicable Decisions

Scenario: A control or obligation is marked not applicable without sufficient rationale.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?

- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Investigate Closed Findings

Scenario: A finding is closed and the auditor must verify closure evidence.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.

- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Investigate Stale Evidence

Scenario: Evidence is old, contextless, generic, or not linked to the claim.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.

- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Investigate Missing Owners

Scenario: Systems, risks, evidence, findings, or remediation lack accountable owners.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.

9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.

- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Investigate Unreviewed Discovery Items

Scenario: Shadow AI signals remain open or stale.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?

- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Investigate Agent Runtime Readiness

Scenario: An agent is marked ready but runtime controls may be incomplete.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.

- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.

- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Gateway Decisions

Scenario: Allowed, blocked, approval-required, or logged decisions must be traceable.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.

6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Review Claw Task Outputs

Scenario: Generated runtime outputs must be reviewed before reliance.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?

- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Prepare Management Questions

Scenario: The auditor needs to ask clear questions without editing records.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.

- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.

- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Document Exceptions

Scenario: The review identifies open issues requiring management response.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.

3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.

- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Assess Reliance

Scenario: The auditor decides whether records are reliable for the intended conclusion.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?

- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Request Updated Workpapers

Scenario: The live workspace changed after a pack or report was generated.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.

- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.
- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Close The Auditor Review

Scenario: The auditor completes review and access should be revoked.

Records and views usually involved:

- Dashboard and workspace overview.
- AI System Records and AI Use Case Intake.
- Risk tiering, ACRS, and routed assessment plan.
- Framework assessment modules.
- Discovery, where enabled.
- Risk Register, Evidence Tracker, Testing Centre, Findings Register, and Remediation Roadmap.
- Model Cards and policy documents.
- Agentic CISO and Gateway records where agents are in scope.
- Audit Trail, Workpaper Packs, Board Dashboard, Reports, History, Baselines, Gap Narratives, Cross-Framework View, and Cert Readiness.

Review procedure:

1. Confirm scope before reviewing details.
2. Identify the source record behind the dashboard, report, workpaper, or claim.
3. Inspect owner, status, date, rationale, evidence, and linked records.
4. Compare the record to the intended conclusion.
5. Sample related records when the claim is material.
6. Note unsupported, stale, inconsistent, or incomplete records.
7. Ask accountable owners to correct records or provide evidence.
8. Recheck updated records before relying on them.
9. Document exceptions, residual uncertainty, and reliance limitations.

Auditor questions:

- What exactly is in scope and out of scope?
- Who owns the system, control, risk, finding, or remediation?
- What evidence proves the claim?
- Is the evidence current and specific?
- Are scores supported by rationale?
- Are not-applicable decisions justified?
- Are high and critical gaps visible in risk, finding, or remediation records?
- Are reports generated from the current workspace?
- Are discovery or agentic runtime claims backed by source records?
- What reliance limitation remains after review?

Evidence to inspect:

- System and intake records.
- Framework scores and notes.
- Evidence records and evidence requests.
- Test methods, samples, results, and retest evidence.
- Findings, management responses, and closure evidence.
- Remediation status, owner, due date, and closure criteria.
- Discovery source, queue, asset, alert, and coverage records.
- Agentic identity, access, tool, approval, data movement, red-team, incident, Gateway, Claw, and workflow records.
- Audit log events for sensitive or late changes.
- Workpaper and report generation date and scope.

Exception indicators:

- The dashboard looks complete but supporting records are thin.
- High scores rely only on policy documents.

- Evidence is generic, stale, unlinked, or ownerless.
- Findings are closed without closure evidence.
- Discovery coverage is low but inventory completeness is claimed.
- Agent runtime readiness is claimed without Gateway, approval, testing, or incident evidence.
- Reports do not match the live workspace.
- Audit trail shows major late changes without explanation.

Completion standard:

- Scope and limitations are clear.
- Material claims have traceable evidence.
- Exceptions are documented for management response.
- Reliance decisions are proportionate to the quality of records.
- Auditor access remains read-only unless explicitly authorised otherwise.

Auditor Reliance Reference

Review area	Strong reliance signal	Weak reliance signal
Scope	Workspace, period, frameworks, and systems are defined.	Scope is implied or inconsistent.
Inventory	Systems and discovered assets are complete and owned.	Material systems are missing or duplicated.
Intake	Use case facts support route and assessment depth.	Intake is vague or incomplete.
Scores	Scores include rationale and evidence.	Scores are unsupported or copied forward.
Evidence	Evidence is current, specific, linked, and reviewed.	Evidence is stale, generic, or missing.
Testing	Operating effectiveness is tested where needed.	Controls are accepted based on design only.
Findings	Issues are clear, owned, and validated at closure.	Findings are vague or closed without proof.
Remediation	Actions are specific, owned, dated, and testable.	Actions are aspirational or overdue.
Discovery	Sources, coverage, queue decisions, and gaps are reviewed.	Shadow AI coverage is unknown.
Agentic	Runtime readiness is supported by identity, access, tools, approvals, tests, and Gateway history.	Agent readiness is asserted without runtime evidence.
Reports	Reports match live records and state limitations.	Reports overstate readiness or use stale records.