Gamut Agentic CISO & Gateway Manual

Customer-ready launch manual

Version	Launch version
Audience	AI governance leads, CISOs, security reviewers, agent owners, risk teams, and authorised runtime operators

Table of Contents

About This Manual

Key Concepts

Recommended Agent Governance Workflow

Access And Roles

Agentic CISO Overview

Agent Register

Creating Governance Actions From An Agent

Organisation Chart

Identity & Ownership

Access Matrix

Tool Permissions

Gateway Tooling

Connector Records

Data Movement

Approval Gates

Agent Risk Tiering

ATF Control Tower

Incident Response

Red-Team Tests

Agentic Evidence Pack

Agentic Board Report

Framework Mapping

Gateway Policy Tests

Using Policy Test Results

Gateway Decision History

Agent Runtime Control

- BYO Agent Control
- Readiness
- Policy Test
- Workflow Planner
- Decision History

Claw Tasks

Claw Task Review

Workflow Studio

Workflow Templates

Workflow Preflight
Approval Queue
External Notifications And Connector Actions
Safe Use Rules
Common Issues
Launch Checklist For Agentic Governance
Comprehensive Feature Handbook
- Agent Register Deep Dive
- Identity And Ownership Deep Dive
- Access Matrix Deep Dive
- Tool Permissions Deep Dive
- Gateway Tooling Deep Dive
- Data Movement Deep Dive
- Approval Gates Deep Dive
- ATF Control Tower Deep Dive
- Incident Response Deep Dive
- Red-Team Testing Deep Dive
- Gateway Runtime Control Deep Dive
- Agentic Runtime Procedure Library
- Agentic Control Reference Tables
- Agentic Launch Readiness Walkthrough

About This Manual

This manual explains how to govern AI agents and agentic workflows in Gamut. It covers Agentic CISO records, Gateway policy tests, runtime readiness, approval gates, tool permissions, Claw task review, and Workflow Studio.

Some features require Advanced or Enterprise access. If a feature is not visible, ask your administrator to confirm your role, workspace access, and subscription plan.

> Important note: Agentic governance decisions can affect people, systems, data, and production operations. Gamut supports review and control, but it does not replace professional judgment, legal advice, regulatory advice, security review, or authorised operational approval.

Key Concepts

Agentic CISO is the governance workspace for AI agents. It helps teams document agent ownership, identity, access, tool permissions, data movement, approval gates, incident response, red-team testing, and readiness.

Gateway is the policy decision layer for agent actions. It evaluates whether a proposed action should be allowed, blocked, or sent for approval based on the saved governance record.

Claw is the governed runtime task layer. Where enabled, it runs bounded tasks through Gamut policy and Gateway controls, then records decisions, outputs, and evidence for review.

Workflow Studio is the visual planning tool for multi-step governed workflows. It helps users design plans, check readiness, and dispatch approved steps under Gateway and Claw controls.

Recommended Agent Governance Workflow

Use this order for a new agent:

1. Register the agent.
2. Assign human, security, technical, and business owners.
3. Record identity and access details.
4. Map tool permissions.
5. Document data flows.
6. Define approval gates.
7. Complete agent risk tiering and ATF maturity review.
8. Add incident playbooks.
9. Run red-team tests.
10. Run Gateway policy tests.
11. Review runtime readiness.
12. Generate the evidence pack and snapshot report.

Do not treat an agent as production-ready until ownership, permissions, approvals, data boundaries, testing, and evidence are complete.

Access And Roles

Agentic CISO and Gateway access depends on your role and plan.

Typical access patterns:

- Viewers and auditors may be able to review agent records but not edit.
- Contributors and assessors may create or update governance records.
- Lead assessors and authorised administrators may approve, dispatch, or manage higher-impact actions where enabled.
- Production runtime actions require elevated permission and the correct subscription tier.

If you can see a record but cannot change it, check edit mode, workspace role, tenant role, and plan entitlement.

Agentic CISO Overview

The Agentic CISO area gives a portfolio view of agents in the selected workspace.

Use it to review:

- Total registered agents.
- High and critical risk agents.
- Agents without named owners.
- Agents without approval gates.
- Tool permission coverage.
- Incident and red-team readiness.
- Evidence pack readiness.
- Gateway policy test results.

The overview is designed to show whether the agent fleet has enough governance coverage to support assurance or production use.

Agent Register

The Agent Register is the starting point for agent governance.

For each agent, record:

- Agent name.
- Identifier.
- Business function.
- Purpose.
- Model provider and model name.
- Orchestration framework.
- Deployment environment.
- Human owner.
- Security owner.
- Technical owner.
- Business owner.
- Agent type.

- Lifecycle status.
- Risk tier.
- Current and target autonomy level.
- Audit logging status.

Use clear names and identifiers. Avoid vague entries such as "AI bot" or "automation agent" unless they are accompanied by a precise purpose and owner.

Creating Governance Actions From An Agent

From an agent record, users may be able to create linked governance work such as:

- A risk in the Risk Register.
- A finding in the Findings Register.
- Evidence requests for key controls.
- ATF control tests.

Use these actions when an agent record reveals an accountability, access, monitoring, approval, or evidence gap.

Organisation Chart

The organisation chart records reporting and authority relationships for agents, owners, and governance roles.

Use it to document:

- Agent owner structure.
- Reporting lines.
- Escalation paths.
- Authority levels.
- Parent and child agent relationships where relevant.

This helps reviewers understand who can approve, suspend, escalate, or remediate an agent.

Identity & Ownership

Identity & Ownership records describe how an agent is represented in systems and access controls.

Record:

- Identity provider.
- Identity type.
- Service account or non-human identity name.
- Credential owner.
- Credential storage approach.
- Authentication method.
- Rotation or review notes.
- Status.

Every production-capable agent should have a clear identity profile. Agents should not act through shared human accounts.

Access Matrix

The Access Matrix documents what an agent can reach and what level of access it has.

Record:

- System name.
- System type.
- Access level.
- Permission scope.
- Data category.
- Risk rating.
- Approval requirement.
- Monitoring requirement.
- Review status.

Use the matrix to distinguish low-risk read access from higher-risk write, delete, external-send, payment, privileged administration, or code-deployment actions.

Tool Permissions

Tool Permissions define the tools an agent can use and the actions it may perform.

Record:

- Tool name.
- Tool provider.
- Tool type.
- Adapter or catalogue entry where applicable.
- Permission level.
- Allowed action categories.
- Data classes.
- Risk rating.
- Approval requirement.
- Audit logging status.
- Runtime enabled status.
- Rate limits or usage constraints.

High-impact capabilities require special care. Treat the following as high-risk or critical until proven otherwise:

- External communication.
- Production data export.
- Financial transactions.

- Permission changes.
- Deletion.
- Code or configuration changes.
- Access to restricted or regulated data.

Tool permission records should be specific. A broad permission such as "use CRM" is weaker than a bounded permission such as "read customer account status for support triage only."

Gateway Tooling

Gateway Tooling shows tool types available for governed use.

Catalogue entries may describe:

- Tool category.
- Supported action types.
- Data classes.
- Risk level.
- Whether approval is required.
- Whether a tenant connection is needed.
- Runtime status.

Use the catalogue to align agent tool permissions with recognised tool types. This makes Gateway decisions more precise.

Gateway Tooling commonly includes four working areas:

- Adapter catalog: The governed catalogue of recognised tool adapters and action categories.
- Tenant connections: Connection records for approved providers or tool backends.
- Runtime parity: Runtime readiness records such as backends, toolsets, skills, checkpoints, and trajectory events.
- Invocation history: Prior Gateway tool invocations, statuses, decisions, reasons, and execution references.

Use Adapter catalog to confirm that a tool permission maps to a known adapter. If an agent has a tool permission but no matching adapter, runtime enforcement may block or warn because the action cannot be evaluated precisely.

Use Tenant connections to confirm whether a tool can actually connect to the approved provider. Check connection name, provider, connection type, owner, workspace scope, allowed actions, allowed data classes, credential status, and last test result. Test a connection after saving it where that action is available.

Use Runtime parity to review:

- Runtime backends: Approved runtime, sandbox, browser, or backend proxy metadata.
- Toolsets: Bundles of adapters and allowed actions assigned to an agent or workflow.
- Skills: Governed reusable capability definitions, including required adapters, input/output expectations, instructions, and policy constraints.
- Checkpoints: Saved state snapshots used to support rollback, review, or repeatability.
- Trajectory events: Runtime or workflow events that show what happened during governed execution.

Use Invocation history to inspect prior tool calls. Review tool name, action type, decision, status, severity, reason, created time, and execution reference. Blocked, failed, rejected, or approval-required invocations should be linked to a control fix, risk, finding, evidence request, approval gate, or incident review where appropriate.

Do not use Gateway Tooling to create broad, vague runtime access. Tool records should describe what the agent can do, which data classes it may touch, whether a tenant connection is required, and whether approval or logging is mandatory.

Connector Records

Where enabled, connector records show tenant or workspace connections used by governed tools.

Use connector records to:

- Confirm which provider or connector is configured.
- Check whether a connector is tenant-wide or scoped to a workspace.
- Validate readiness before runtime use.
- Revoke or retest connections where permitted.

Do not store credentials in notes, descriptions, prompts, or workflow objectives. Use the approved connector configuration flow.

When reviewing connector records, check:

1. The connector is owned by an accountable person or team.
2. The connector is scoped correctly: tenant-wide only when genuinely required, workspace-specific when possible.
3. Credential status is healthy or intentionally not configured.
4. Allowed actions and data classes match the governance need.
5. Last test status is recent enough for the planned use.
6. The connector is not being used to bypass an approval gate or data movement restriction.

Data Movement

Data Movement records describe how data moves through the agent or workflow.

Record:

- Source system.
- Destination system.
- Data type.
- Data classification.
- Movement type.
- Whether export is allowed.
- Encryption or protection requirements.
- Retention notes.
- Risk rating.

Data movement records help reviewers understand whether an agent can move personal, confidential, restricted, or critical data outside the intended boundary.

Approval Gates

Approval Gates define when human approval is required.

Record:

- Gate name.
- Agent.
- Action category.
- Tool or target system.
- Data class.
- Environment.
- Risk level.
- Approval method.
- Approver role or named approver.
- Evidence required.
- Status.

Approval gates should be active, specific, and aligned to the action risk. A broad approval rule may be useful as a backstop, but high-impact actions should have precise gates.

Examples of actions that usually need approval:

- External customer communication.
- Financial transactions.
- Production data export.
- High-risk code or configuration changes.
- Restricted data access.
- Critical incident actions.

Agent Risk Tiering

Agent risk tiering helps classify the inherent risk of an agent.

Consider:

- Autonomy level.
- Data access.
- External impact.
- Tool permissions.
- Human oversight.
- Business criticality.
- Regulatory or safety exposure.

- Whether actions are reversible.

Risk tiering should drive the depth of controls, approval gates, monitoring, red-team testing, and evidence required.

ATF Control Tower

ATF Control Tower summarises agent maturity and readiness against the Agentic Trust Framework.

Use it to:

- Compare current and target autonomy.
- Review whether required controls exist.
- Identify missing owner or status evidence.
- Create ATF control tests.
- Link back to the framework assessment where needed.

Do not promote an agent to a higher autonomy level until required ownership, approval, access, logging, testing, and incident response controls are in place.

Incident Response

Incident Response records document what to do when an agent fails, behaves unexpectedly, or is suspected of being compromised.

Record:

- Incident type.
- Severity.
- Detection trigger.
- Initial response.
- Containment steps.
- Escalation path.
- Evidence preservation.
- Recovery steps.
- Communication plan.
- Post-incident review.
- Test status.

Important playbook types include prompt injection compromise, unauthorised external communication, unsafe delegation, model or provider compromise, excessive autonomy, and policy bypass attempts.

Red-Team Tests

Red-Team Tests help prove whether agentic controls work under realistic challenge conditions.

Record:

- Test name.

- Test type.
- Scenario.
- Objective.
- Agent.
- Expected control.
- Result.
- Severity.
- Findings or remediation links.
- Tester and test date.

Failed red-team tests should usually become findings or remediation actions. Retest after remediation.

Agentic Evidence Pack

The Agentic Evidence Pack summarises the governance evidence for the agent fleet.

It may include:

- Registered agents.
- Ownership status.
- Identity profiles.
- Access matrix coverage.
- Tool permissions.
- Approval gates.
- Data flows.
- Incident playbooks.
- Red-team tests.
- Gateway policy decisions.
- ATF readiness.

Use the evidence pack for audit preparation, management review, client assurance, and readiness checks.

Agentic Board Report

Agentic Board Report provides a board-ready summary of agentic governance posture.

It typically highlights:

- Number of agents in scope.
- High and critical risk agents.
- Missing owners.
- Missing approval gates.
- Red-team status.
- Tool and access risks.

- Board attention items.
- Recommended actions.

Use it for executive reporting, not as a substitute for detailed evidence review.

Framework Mapping

Framework Mapping shows how agentic governance domains align with broader frameworks.

Domains may include:

- Identity and ownership.
- Tool permissions.
- Human approval gates.
- Data movement.
- Kill switch and incident response.
- Red-team testing.

Use the map to understand where one agentic control supports multiple assurance expectations.

Additional agentic report views may include Security Architecture Report, Red Team Report, and Control Tower Report. Use these reports to package focused slices of the same underlying governance record for the right audience.

Gateway Policy Tests

Gateway policy tests evaluate a proposed agent action against the saved governance record.

A policy test typically asks:

- Which agent is acting?
- What tool is being used?
- What action is requested?
- What target system is involved?
- What data classification applies?
- Is the environment production or non-production?
- What is the external impact?
- Is approval required?

The result may be:

- Allow.
- Allow with logging.
- Require review.
- Require approval.
- Block.
- Block and create finding.
- Block and open incident.

Treat a Gateway result as a governance decision aid. If the decision is blocked or requires approval, resolve the underlying control gap rather than bypassing the result.

Using Policy Test Results

After a policy test, review:

- Decision.
- Severity.
- Reason.
- Controls passed.
- Controls failed.
- Missing controls.
- Missing evidence.
- Required approval.
- Recommended action.
- Policy snapshot.

Where available, create linked records directly from the result:

- Finding.
- Risk.
- Evidence request.
- Approval gate.

This keeps policy decisions connected to accountable remediation.

Gateway Decision History

Gateway Decision History records prior policy tests and runtime decisions.

Use history to:

- Review blocked or approved actions.
- Re-run a prior scenario.
- Ask the AI Consultant about a decision.
- Delete test records where permitted.
- Package decisions into evidence or reporting.

History is useful for showing how governance improved over time.

Agent Runtime Control

Agent Runtime Control evaluates whether an agent has enough governance coverage for runtime use.

Readiness may consider:

- Agent registration.

- Named owners.
- Tool permissions.
- Tool catalogue mapping.
- Approval gates.
- Data movement records.
- Prior policy tests.
- Audit logging.
- Connector readiness.
- Red-team and incident readiness.

An agent marked not ready should not be treated as safe for unrestricted runtime use. Resolve readiness blockers first.

The Agent Runtime Control Centre may include these tabs:

- BYO Agent Control: Enterprise runtime control for externally hosted agents that act through Gamut-governed credentials and Gateway policy.
- Readiness: Per-agent readiness calculated from Agent Register, Tool Permissions, Approval Gates, Data Movement records, and prior Gateway decisions.
- Policy test: A simulator for proposed agent actions.
- Workflow planner: Governed multi-step workflow planning and dispatch.
- Claw tasks: Governed runtime task execution and review.
- Decision history: Saved policy tests and runtime decisions.

BYO Agent Control

BYO Agent Control is for agents that run outside Gamut but must act through Gamut governance. A registered agent becomes a BYO runtime agent only when an authorised runtime credential is issued or an external heartbeat is observed.

Use BYO Agent Control to review:

- Total, issued, healthy, and blocked BYO credentials.
- Active BYO agents and child runtimes.
- Unhealthy credentials and stale heartbeats.
- Pending approvals.
- Blocked runtime actions.
- Governance parity blockers by category.
- External agent events.
- Zero-trust controls such as Gateway-only execution, credential health, replay protection, tenant binding, and fail-closed handling.

Before issuing or relying on BYO runtime credentials, confirm:

1. The agent is registered and has a stable identifier.
2. Human, security, technical, or business ownership is recorded.

3. Lifecycle status and risk tier are current.
4. ATF level is assigned.
5. Tool permissions include Gateway adapter keys, runtime-enabled status, audit logging, allowed actions, and allowed data classes.
6. Required approval gates are active.
7. Data movement records cover external or high-impact actions.
8. Incident response and red-team evidence are adequate for the risk tier.

Do not issue runtime credentials to compensate for incomplete governance records. Fix the parity blocker first.

Readiness

The Readiness tab evaluates whether each agent can safely proceed toward runtime use.

Readiness checks may flag:

- Missing agent identifier.
- Missing named owner.
- Missing ATF level.
- Tool-capable agent with no active tool permission.
- Tool permission without Gateway adapter key.
- Audit logging disabled.
- Runtime not enabled.
- High-risk tool without an approval gate.
- External or high-impact tool without a data movement record.
- No policy test history.
- Tool requiring a tenant connection that is not configured.

Use Readiness as a pre-production gate. An agent can be documented but still not be runtime-ready.

Policy Test

The Policy test tab evaluates a proposed action before it runs.

When creating a policy test, define:

- Agent or scenario template.
- Tool name.
- Requested action or action type.
- Target system.
- Data classification.
- Environment.
- External impact.
- Human requestor.
- Workflow name or purpose.

After running the test, review the decision, severity, controls passed, controls failed, missing controls, missing evidence, approval requirement, and policy snapshot. Use the linked actions to create findings, risks, evidence requests, or approval gates from material gaps.

High-risk agents, production workflows, and critical tools should have recent policy tests. Stale policy tests should be repeated after material changes.

Workflow Planner

The Workflow planner tab supports governed multi-step runtime plans. It is distinct from static documentation because steps may be dispatched under Gateway and Claw controls where permitted.

Use it to:

- Create or load a workflow plan.
- Add steps and dependencies.
- Select task types, agents, tools, objectives, and required approvals.
- Review preflight status.
- Approve and dispatch where authorised.
- Sync runtime state.

Do not dispatch a workflow until preflight blockers are resolved and required approvals are complete.

Decision History

Decision History stores prior policy tests and runtime decisions. Use it to:

- Review allowed, blocked, approval-required, and logged decisions.
- Re-run an old scenario after controls change.
- Ask the AI Consultant to review a decision.
- Delete test records where permitted.
- Package relevant decisions into evidence or reporting.

Decision history is valuable only when users act on blocked or stale decisions. A blocked decision without a linked action is an unresolved governance issue.

Claw Tasks

Where enabled, Claw tasks run bounded governed work under Gamut policy.

Task types may include:

- Governed context summary.
- Evidence gap analysis.
- Control recommendation.
- Investigation.
- Evidence collection.
- Incident triage.
- Report drafting.
- Control testing.

- Risk review.
- Gateway tool task.

Before a task runs, Gamut checks preflight conditions. If required controls are missing, the task is blocked.

Claw Task Review

Use Claw task history to review:

- Task type.
- Status.
- Agent.
- Gateway decisions.
- Bounded output preview.
- Errors or approval requirements.
- Runtime journal events.
- Materialised outputs.

Where enabled, users can:

- Open a stream.
- Retry a failed or approval-required task.
- Close a stream.
- Materialise output into evidence notes, findings, or control recommendations.
- Pause or resume schedules.

Always review generated outputs before relying on them. Materialised records should be validated by an accountable owner.

Workflow Studio

Workflow Studio is a visual planner for governed multi-step workflows.

Use it to:

- Select a workspace.
- Add task nodes.
- Apply templates.
- Create dependencies between steps.
- Require approval for specific steps.
- Save governed plans.
- Sync runtime state.
- Review preflight blockers.
- Dispatch approved steps where permitted.

Workflow Studio does not bypass Gamut governance. Gateway and Claw still enforce runtime controls.

Workflow Templates

Workflow Studio includes templates such as:

- Board briefing.
- Risk digest.
- Control status summary.
- Assessment progress snapshot.
- Evidence gap summary.
- Audit trail summary.
- Evidence pack refresh.
- Approval queue reminder.
- Overdue action escalation.
- Control owner follow-up.
- External notification.
- Secure connector task.
- Agent risk review.
- Security incident triage.
- Ephemeral research branch.
- Marketing campaign guardrail.

Templates are starting points. Review each step, dependency, tool, approval requirement, and output destination before saving or dispatching.

Workflow Preflight

Preflight checks whether a workflow is ready.

Common blockers include:

- No registered agents.
- Missing owners.
- Missing tool permissions.
- Tool logging disabled.
- Missing approval gates.
- Missing connector readiness.
- Runtime infrastructure unavailable.
- Cyclic workflow dependencies.
- Required approval not configured.

Use the suggested fix actions to return to the relevant Agentic CISO screen and resolve blockers.

Approval Queue

Some Gateway or workflow actions require approval before continuing.

Good approval practice:

- The requester should not approve their own high-impact action.
- The approver should review the action, data class, target system, expected output, and risk.
- The approver should provide a written justification.
- Rejected approvals should result in a finding, risk, or remediation action where appropriate.

Approvals are governance records, not informal permission slips.

External Notifications And Connector Actions

External notifications and connector actions require extra caution because they can affect people, systems, customers, or third parties.

Before approving or dispatching:

- Confirm the message or action is bounded.
- Confirm the tool is approved.
- Confirm the target system is correct.
- Confirm the data class is appropriate.
- Confirm evidence supports the action.
- Confirm logging and approval are in place.
- Confirm external sharing or system action has been approved by the accountable owner.

Do not use workflow text to store secrets, credentials, private keys, or unrestricted sensitive data.

Safe Use Rules

Follow these rules for agent governance:

- Every agent must have a named human owner.
- Every high-risk agent should have a named security owner.
- Every production-capable agent should have identity and access records.
- Every critical tool permission should require approval and audit logging.
- Every external, financial, data-export, delete, code-change, or privileged action should be reviewed carefully.
- Every failed red-team test should be tracked to closure.
- Every blocked Gateway decision should lead to a control fix, finding, risk, or evidence request.
- Every generated output should be reviewed before use.

Common Issues

I cannot see Agentic CISO:

- Your role or plan may not include it.

- Ask your administrator to review access.

I can see Agentic CISO but cannot edit:

- Unlock edit mode.
- Check your workspace role.
- Confirm the workspace is not read-only.

Gateway policy test is blocked:

- Review controls failed and missing controls.
- Add or update tool permissions.
- Add approval gates.
- Add missing owners or data flow records.
- Re-run the policy test after updates.

Claw task is blocked by preflight:

- Review the missing readiness items.
- Confirm the agent, tool permissions, approval gates, data flows, and connector readiness.
- Resolve blockers before retrying.

Workflow cannot dispatch:

- Save the plan.
- Review preflight.
- Resolve missing agents, tools, approvals, dependencies, or runtime readiness.
- Approve the plan where required.

Output should become a finding or evidence note:

- Use the available materialise action.
- Review and edit the created record.
- Assign an owner and due date where relevant.

Launch Checklist For Agentic Governance

Before using an agent in production:

- Agent is registered.
- Human owner is assigned.
- Security owner is assigned.
- Lifecycle status is current.
- Identity profile is documented.
- Access matrix is complete.
- Tool permissions are mapped.
- Critical tools require approval.

- Audit logging is enabled for relevant tools.
- Data flows are documented.
- Approval gates are active and specific.
- Risk tier is confirmed.
- ATF current and target levels are recorded.
- Incident playbooks exist and have been tested.
- Red-team testing has been completed.
- Failed tests have findings or remediation actions.
- Gateway policy tests have been run for high-impact actions.
- Runtime readiness has no unresolved blockers.
- Evidence pack and snapshot report have been reviewed.

Comprehensive Feature Handbook

This handbook layer is for teams governing agentic and runtime AI. Treat it as a control reference before any agent is allowed to act in production or touch high-impact tools.

Agent Register Deep Dive

Document every agent's purpose, owners, model, framework, environment, lifecycle, autonomy, capabilities, risk, approval status, and review dates.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.

- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Identity And Ownership Deep Dive

Define non-human identity, credential owner, authentication method, storage approach, review frequency, and approval requirements.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Access Matrix Deep Dive

Record systems, access levels, data classes, action types, approval needs, logging, monitoring, and review frequency.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.

2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Tool Permissions Deep Dive

Map tools, adapters, allowed actions, data classes, runtime status, approval requirements, limits, sandboxing, and logging.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Gateway Tooling Deep Dive

Review adapter catalog, tenant connections, runtime parity records, toolsets, skills, checkpoints, trajectory events, and invocation history.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.

7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Data Movement Deep Dive

Document source, destination, data type, classification, personal data flags, retention, masking, encryption, export allowance, and risk.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Approval Gates Deep Dive

Define when approval is needed, who approves, what evidence is required, how timeout is handled, and escalation path.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.

- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

ATF Control Tower Deep Dive

Use current and target maturity levels to manage autonomy promotion and evidence requirements.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.

- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Incident Response Deep Dive

Prepare detection, containment, suspension, access revocation, evidence preservation, communication, recovery, and review procedures.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Red-Team Testing Deep Dive

Test prompt injection, unsafe delegation, tool misuse, privilege escalation, data leakage, policy bypass, and recovery controls.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.

- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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.

Gateway Runtime Control Deep Dive

Use BYO Agent Control, readiness, policy tests, workflow planner, Claw tasks, and decision history before production actions.

Primary users: agent owners, security reviewers, governance leads, and authorised runtime operators.

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:

- Agent: The registered agent or workflow subject.
- Owner: The accountable human, security, technical, or business owner.
- Risk level: The severity or risk tier driving control depth.
- Approval requirement: Whether human approval is needed before action.
- Evidence: Proof that the control exists and operates.
- Status: Draft, active, blocked, retired, tested, failed, or equivalent state.

How to use this feature:

1. Start from the Agent Register and confirm the agent is in scope.
2. Complete owner, identity, access, tool, data movement, and approval records.
3. Run ATF and risk review before enabling high-impact capability.
4. Create incident and red-team records before production use.
5. Run Gateway policy tests for material action types.
6. Resolve readiness blockers before dispatch, credential issue, or Claw task use.
7. Review decision history and runtime evidence after execution.

Quality checks before you move on:

- No production-capable agent lacks a human owner.
- No high-impact tool lacks approval and logging.
- No external data movement lacks a data flow record.
- No failed red-team test is left without action.
- No blocked Gateway decision is ignored.

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 Runtime Procedure Library

The following procedures are designed for launch teams operating agentic governance, Gateway controls, Claw tasks, BYO agents, and workflow dispatch. They are intentionally detailed because runtime actions can affect systems, people, data, and external parties.

Register A New Agent

Scenario: A team has created or purchased an agent and needs to place it under governance.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.

3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.

- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Classify Agent Autonomy

Scenario: The owner needs to decide whether the agent is Intern, Junior, Senior, or Principal level.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.

- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?

- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.

- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Assign Agent Ownership

Scenario: The agent lacks named human, business, security, or technical owners.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.

6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create An Identity Profile

Scenario: The agent will authenticate to systems or act through a non-human identity.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.

- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?

- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.

- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Review Service Account Access

Scenario: The agent uses a service account or other non-human credential.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.

10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.

- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Build An Access Matrix

Scenario: The team needs to document which systems and data the agent can reach.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.

- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Add Tool Permissions

Scenario: The agent needs to call tools, APIs, connectors, or runtime adapters.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Map Tool Permissions To Gateway Adapters

Scenario: A tool permission exists but Gateway needs a precise adapter mapping.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.

3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.

- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Configure Tenant Tool Connections

Scenario: A governed tool requires an approved tenant or workspace connection.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.

- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?

- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.

- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Define A Runtime Backend

Scenario: The organisation needs a governed runtime, sandbox, browser, or backend proxy record.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.

6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create A Toolset

Scenario: Several adapters and allowed actions must be bundled for an agent or workflow.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.

- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?

- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.

- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create A Skill

Scenario: A reusable governed capability needs instructions, required adapters, and input/output expectations.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.

10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.

- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Record A Checkpoint

Scenario: A workflow or task needs a reviewable state snapshot for rollback or assurance.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.

- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Record A Trajectory Event

Scenario: A runtime event needs to be preserved as part of traceability.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Document Data Movement

Scenario: The agent moves data between systems, environments, or external destinations.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.

3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.

- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create An Approval Gate

Scenario: A high-impact action needs human approval before execution.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.

- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?

- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.

- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Review Pending Gateway Approvals

Scenario: An action is waiting for authorised approval or rejection.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.

6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Run A Gateway Policy Test

Scenario: The team wants to evaluate a proposed action before execution.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.

- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?

- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.

- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Interpret A Blocked Gateway Decision

Scenario: Gateway blocks an action or requires approval because controls are missing.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.

10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.

- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create A Finding From A Gateway Gap

Scenario: A failed policy test reveals a control weakness.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.

- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create A Risk From A Gateway Gap

Scenario: A failed policy test reveals exposure that may continue if untreated.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create An Evidence Request From A Gateway Gap

Scenario: A policy test cannot pass because required proof is missing.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.

3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.

- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Review Gateway Decision History

Scenario: The team needs to understand prior policy decisions and changes over time.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.

- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?

- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.

- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Enable BYO Agent Control

Scenario: An externally hosted agent must act through governed credentials and Gateway policy.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.

6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Review BYO Governance Parity

Scenario: A BYO agent has blockers across identity, ownership, risk, tools, approvals, data movement, red team, or incident response.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.

- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?

- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.

- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Investigate Stale Heartbeats

Scenario: A BYO runtime has not reported activity within the expected period.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.

10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.

- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Respond To Unhealthy Credentials

Scenario: A runtime credential is unhealthy, missing, expired, or blocked.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.

- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Register A Child Runtime

Scenario: A delegated child agent needs governance parity under a parent agent.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Run A Claw Preflight

Scenario: A Claw task must prove that required tools, adapters, and approvals are ready.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.

3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.

- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Run A Claw Task

Scenario: A governed runtime task is ready to execute within policy constraints.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.

- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?

- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.

- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Review A Claw Task Result

Scenario: A completed task has output, journal events, or Gateway decisions that require review.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.

6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Retry A Failed Claw Task

Scenario: A task failed or required approval and the team wants to retry after fixing blockers.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.

- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?

- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.

- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Materialise Claw Output

Scenario: Reviewed output should become an evidence note, finding, or control recommendation.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.

10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.

- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create A Claw Schedule

Scenario: A recurring governed task is needed for periodic review or monitoring.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.

- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Pause Or Resume A Claw Schedule

Scenario: A scheduled task must be stopped or restarted safely.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Plan A Governed Workflow

Scenario: A multi-step workflow needs steps, dependencies, agents, task types, and approvals.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.

3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.

- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Use A Workflow Template

Scenario: A user wants to start from a board briefing, risk digest, evidence gap summary, incident triage, or connector task template.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.

- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?

- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.

- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Resolve Workflow Preflight Blockers

Scenario: A workflow cannot dispatch because a dependency, agent, tool, approval, or runtime control is missing.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.

5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.

- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Dispatch A Workflow

Scenario: An approved workflow is ready for Gateway and Claw controlled execution.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.

- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?

- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.

- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Review Workflow Runtime State

Scenario: A dispatched workflow has step results, approvals, or execution references to inspect.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.

8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.

- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create An Incident Playbook

Scenario: The agent needs a response plan for compromise, unsafe delegation, policy bypass, external communication, or excessive autonomy.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.

- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.

- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Test An Incident Playbook

Scenario: The team needs evidence that response procedures work.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.

11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.

- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Create A Red-Team Test

Scenario: The team needs to challenge agent controls with realistic adversarial scenarios.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.

- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Respond To A Failed Red-Team Test

Scenario: A red-team test fails and must become a finding or remediation action.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.

- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Promote ATF Maturity

Scenario: The owner wants to increase current or target autonomy level.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.

4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.

- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Demote Or Suspend An Agent

Scenario: The agent no longer meets governance, evidence, risk, or runtime expectations.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.

- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?

- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Generate An Agentic Evidence Pack

Scenario: The team needs evidence for audit, client assurance, or readiness review.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.

8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.

- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Prepare An Agentic Board Report

Scenario: Leadership needs a concise agentic governance posture update.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.

- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.

- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Use The AI Consultant For Gateway Review

Scenario: A user needs help explaining a policy decision and remediation path.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.

11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.

- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Run A Quarterly Agent Review

Scenario: The organisation needs periodic review of owners, access, approvals, incidents, tests, and readiness.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.

- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.
- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Retire An Agent

Scenario: An agent is no longer used and must be closed without losing auditability.

Records and views usually involved:

- Agent Register.
- Organisation Chart where reporting, parent, child, or escalation relationships matter.
- Identity & Ownership.
- Access Matrix.
- Tool Permissions.
- Gateway Tooling, including adapter catalog and tenant connections.
- Data Movement.
- Approval Gates.
- ATF Control Tower.
- Incident Response.
- Red-Team Tests.
- Gateway Policy Test, Readiness, BYO Agent Control, Claw Tasks, Workflow Planner, or Decision History where runtime is involved.
- Risk Register, Evidence Tracker, Findings Register, Testing Centre, and Remediation Roadmap for follow-up.

Procedure:

1. Confirm the workspace and the agent or workflow in scope.
2. Open the Agent Register and verify the agent name, identifier, purpose, lifecycle status, environment, model/provider context, owners, autonomy level, capability flags, risk tier, approval status, and review dates.
3. Check Identity & Ownership. A production-capable agent should not act through a shared human account or undocumented credential path.
4. Review the Access Matrix for every system, data category, action type, approval requirement, logging requirement, monitoring requirement, and review frequency relevant to the scenario.
5. Review Tool Permissions. Confirm adapter key, allowed data classes, runtime enabled status, action permissions, transaction limits, approval requirement, sandboxing, audit logging, risk rating, and status.
6. If the tool requires a connector, open Gateway Tooling and confirm tenant connection readiness, owner, scope, allowed actions, allowed data classes, credential status, and last test result.
7. Review Data Movement where information leaves a system, crosses environments, touches customer or personal data, is retained, or may be exported.
8. Review Approval Gates for high-impact, external, financial, deletion, permission-change, production, restricted-data, or code/configuration actions.
9. Run or review Gateway readiness and policy tests before any live runtime action.
10. If Gateway blocks or requires approval, resolve the control gap through records, evidence, approvals, risks, findings, or remediation rather than bypassing the decision.
11. For Claw or workflow execution, complete preflight and confirm required adapters, audit logging, approvals, and runtime readiness.
12. After runtime activity, review decision history, task output, journal events, materialised records, and evidence links.
13. Update review dates and next actions so the agent record remains live.

Control questions:

- Is the agent's purpose specific enough to test and govern?
- Is there a named human owner accountable for outcomes?
- Is there a security owner for high-risk or production-capable agents?
- Does the identity model distinguish the agent from human users?
- Are tool permissions narrower than the broad system permissions available in the connected platform?
- Are data classes and action types explicit?
- Are high-impact actions blocked or approval-gated?
- Does Gateway have enough evidence to allow the action?
- Does Claw or workflow execution preserve a reviewable trail?
- Would an auditor understand why the action was allowed, blocked, or approved?

Evidence to collect:

- Agent design or ownership documentation.
- Service account or non-human identity approval.
- Access review or permission matrix.
- Tool catalogue, adapter, and connector readiness records.
- Approval gate configuration and approval history.
- Data movement assessment and retention/protection evidence.
- Gateway policy test result and decision path.
- Claw task journal, bounded output summary, and materialised records where used.
- Incident playbook and test evidence.
- Red-team test scenario, result, finding, and remediation evidence.
- ATF maturity rationale and promotion or demotion decision.

Decision points:

- If the agent lacks an owner, do not treat it as production-ready.
- If the identity path is shared or unclear, create an identity gap before enabling runtime use.
- If a tool can write, delete, send externally, spend money, change permissions, export data, or modify code, require stronger review.
- If data movement is not documented, pause external or high-impact actions.
- If no approval gate exists for a high-risk action, create one before execution.
- If Gateway blocks, use the blocked decision as a control signal.
- If Claw preflight fails, fix governance records before retrying.
- If a red-team test fails, create a finding or remediation and retest before promotion.
- If readiness is stale, rerun policy tests after material changes.

Completion standard:

- Agent, identity, access, tools, data movement, approval, incident, test, and runtime records are complete enough for the risk tier.

- The action or runtime decision can be traced to saved governance facts.
- Missing controls have linked risks, findings, evidence requests, or remediation.
- High-impact runtime actions have approval and audit history.
- The next review date and owner are clear.

Failure indicators:

- Agent listed as active but no human owner is recorded.
- Tool permission exists but no adapter key or connector readiness exists.
- Approval gate is broad, inactive, or not linked to the high-impact action.
- Data movement says export is allowed but no retention, masking, or encryption expectation is recorded.
- Gateway decision history contains blocked tests without linked action.
- Claw output is used directly without reviewer validation.
- ATF maturity is raised without evidence period, red-team results, or incident readiness.

Agentic Control Reference Tables

Control area	What good looks like	Primary Gamut evidence
Identity	Agent has a unique identity, credential owner, authentication method, review cadence, and expiry controls.	Identity & Ownership, Access Matrix, audit logs
Ownership	Human, business, security, and technical ownership are recorded for the agent's risk level.	Agent Register, Organisation Chart
Least privilege	Access and tools are limited to the actions and data classes required for the purpose.	Access Matrix, Tool Permissions, Gateway Tooling
Approval	High-impact actions require approval by someone other than the requester where appropriate.	Approval Gates, Gateway approvals, Decision History
Data boundary	Data flows, classification, retention, masking, encryption, and export rules are documented.	Data Movement
Monitoring	Actions are logged, decisions are traceable, and anomalies can be investigated.	Gateway Decision History, Claw journal, audit logs
Incident readiness	Containment, suspension, access revocation, evidence preservation, communication, recovery, and review steps exist.	Incident Response
Adversarial testing	Prompt injection, unsafe delegation, tool abuse, data leakage, policy bypass, and privilege escalation have been tested.	Red-Team Tests
Runtime readiness	Gateway, Claw, BYO credentials, workflow dispatch, and connectors are ready only after preflight passes.	Agent Runtime Control Centre
Assurance packaging	Evidence pack and board report reflect actual records and unresolved gaps.	Evidence Pack, Board Report

Agentic Launch Readiness Walkthrough

Use this walkthrough before an agent is approved for production or high-impact runtime use.

1. Open the Agent Register and confirm the agent is active only if the lifecycle state is accurate.
2. Confirm the agent identifier is stable and unique.
3. Confirm human owner, security owner, technical owner, and business owner where relevant.

4. Confirm capability flags: delegation, memory, tool calls, external action, spend authority, code changes, customer data, and confidential data.
5. Confirm current and target autonomy levels.
6. Confirm risk tier and approval status.
7. Open Identity & Ownership and verify credential owner, storage approach, authentication method, access expiry, and review frequency.
8. Open Access Matrix and verify each system, data class, access level, action type, approval requirement, logging requirement, and monitoring requirement.
9. Open Tool Permissions and verify adapter key, allowed data classes, runtime enabled status, action flags, limits, sandbox status, approval requirement, audit logging, risk rating, and status.
10. Open Gateway Tooling and verify adapter catalog mapping, tenant connection readiness, runtime backend, toolset, skill, checkpoints, trajectory evidence, and invocation history where relevant.
11. Open Data Movement and verify source, destination, classification, personal data flags, retention, masking, encryption, export allowance, and risk rating.
12. Open Approval Gates and verify that every high-impact action has an active and specific gate.
13. Open Incident Response and verify containment, suspension, access revocation, evidence preservation, communication, regulator consideration, recovery, post-incident review, owner, and test status.
14. Open Red-Team Tests and verify that realistic tests have been completed, failed tests have findings/remediation, and retests are planned or completed.
15. Open ATF Control Tower and confirm the current level is justified by evidence.
16. Open Agent Runtime Control Centre and review readiness.
17. Run Gateway policy tests for high-impact actions.
18. Resolve blocked tests and missing evidence.
19. If using BYO Agent Control, verify credential health, heartbeat recency, governance parity, child runtimes, pending approvals, and external events.
20. If using Claw or Workflow Planner, run preflight before execution and review output after execution.
21. Generate the Evidence Pack and Board Report only after unresolved blockers are clearly documented.

An agent is not launch-ready simply because the register record exists. It is launch-ready only when the runtime path, approval model, evidence, monitoring, incident response, and testing are all proportionate to the risk tier.