ODA3 INSTITUTE

EXECUTIVE / PRACTITIONER GUIDE

The CISO's Guide to AI Security Assurance

A Decision Guide to Controls, Evidence, Incident Readiness and Executive Accountability

Doc ID: ODA3-2026-08-WHP-COM-018
Classification: Public — GEL v1.0

The CISO's Guide to AI Security Assurance

Table of Contents

ODA3 INSTITUTE

THE CISO'S GUIDE TO AI SECURITY ASSURANCE

A Decision Guide to Controls, Evidence, Incident Readiness and Executive Accountability

Doc ID: ODA3-2026-08-WHP-COM-018
Document Type: WHP — Executive / Practitioner Decision Guide
Framework Context: GAISSF Ecosystem / Operational AI Assurance
Classification: Public — GEL v1.0
Status: CORRECTED FINAL (R10)
Date: 19 August 2026 (supersedes R9, 13 August 2026)
Publisher: ODA3 Institute

Primary audience: CISOs, Deputy CISOs, AI Security Leads, Security Architecture Leaders, Enterprise Security Leaders, SOC/CSIRT Leadership, AI Governance Security Leads and Enterprise Risk/Security Leadership.
Secondary audience: CIO/CTO functions, system/service owners, engineering/platform leaders, GRC, procurement/supplier risk, internal assurance, legal/privacy, safety/OT leadership and authorized delivery partners where relevant.

Publication and Authority Boundary

This is an informative CISO decision guide. It does not create, amend or replace GAISSF™ normative requirements, UAIF™ classification rules, AI-IRF™ response requirements, PAI-SF™ physical-AI controls, controlled assessment methodology, certification logic, law, regulation, functional safety, product-safety requirements or organizational risk authority.

Binding framework requirements come only from the current authoritative sources; this guide does not create them. Detailed implementation belongs in **COM-017 — Implementing GAISSF™**. Evidence-pack preparation belongs in **COM-015 — Building an AI Assurance Evidence Pack**. The public practitioner journey ends at evidence preparation; any later assessment activity is separately governed and authorized. Active incident classification and

response belong in **COM-011 / COM-012 / COM-013**, UAIF™ and AI-IRF™. Detailed physical-AI treatment belongs in **COM-014 / PAI-SF™**.

The following source-state summary governs the terminology and boundaries used in this guide. Exact requirements, applicability and licensing remain controlled by the current authoritative sources.

Current Source and Framework State

Topic	Public treatment used in COM-018
Market-facing GAISSF name	**GAISSF™ — Global AI Safety and Security Framework.** Historical source titles may preserve earlier expansion or mark treatment for source identification only.
GAISSF architecture	**59 controls across nine domains.**
Foundational / D9 treatment	**52 canonical controls in D1–D8** form the Foundational baseline. Under the verified controlling source state, the seven D9 controls are additional mandatory controls whenever physical AI or cyber-physical actuation is within the applicable/assessed scope, and D9 applicability is to be considered and justified. This is a source-derived requirement, not a COM-018 rule.
GAISSF source precedence	GAISSF-NOR-001 remains the primary framework authority; NOR-004 is the canonical control library; current release/reconciliation records govern supersession and publication-state interpretation. Historical conflicting wording is not treated as current public authority.
UAIF™ public state	**Final Publication v1.0.**
UAIF machine-readable asset	**TCR-STD-003 / the UAIF machine-readable schema package already exists and is published as the machine-readable companion.** COM-018 therefore describes the schema as published rather than as merely future development.
AI-IRF™ technical architecture	**Final Publication v1.0 — 79 controls / 19 domains.** Historical pre-publication descriptors are not carried forward.
AI-IRF machine-readable asset	**AIIS v1.0.0 machine-readable schema is published** for control/playbook encoding. Its optional `conformance_test` structure is reserved for future population.
PAI-SF™	Current canonical public structure: **12 domains / 41 controls**. GAISSF D9 is retained for continuity; PAI-SF™ supersedes and expands the detailed physical-AI security-assurance layer represented by D9 without creating a second independent obligation.

Topic	Public treatment used in COM-018
Safety boundary	PAI-SF™ interfaces with functional safety and sector assurance but does not replace them or create product/sector approval or physical return-to-service authority.
GEL v1.0	Governs Internal Use, Commercial Use, factual-alignment claims, marks and reserved certification/audit authority. Internal Use rights do not automatically authorize a third party's commercial delivery.
Commercial / programme status	Licensing, partner/empanelment status, assessment authority and certification authority are separate legal/programme authorizations. COM-018 creates none of them and does not state that any such service or status is currently available.

Machine-readable terminology rule for COM-018: when discussing current ODA3 technical assets, use **"published machine-readable schema"** where a schema already exists. Reserve **"development"** for future interoperability work such as product integrations, field mappings, connectors, SDKs, adapters, test harnesses or additional machine-readable profiles—not for the already-published core schema itself.

Guidance and Claims Language

COM-018 is informative practitioner guidance.

- `SHALL` / `SHALL NOT` are not used to invent practitioner obligations; binding obligations come from controlling sources.
- `must` / `mandatory` are reserved for accurate restatement of a verified controlling requirement or an unavoidable source/authority condition.
- `should`, `may`, `consider`, `useful` and similar language describe practitioner guidance rather than new GAISSF requirements.
- A local portfolio label, dashboard status, implementation label, exception status or operating decision is not an ODA3 assessment state, control determination, readiness designation or certification result.
- A checklist, charter, portfolio register or decision playbook in COM-018 is a management aid, not an assessor workpaper, mandatory gate, scoring input, certification precursor or ODA3 approval instrument.
- A CISO management assurance position is not an independent assurance opinion or a controlled GAISSF assessment determination.

Practitioner Suite Role Register

The following practitioner-suite role boundaries govern how COM-018 should be read alongside adjacent ODA3 publications:

Publication	Role retained in COM-018	Consistency rule
COM-008 — The AI Assurance Imperative: A Board Member's Guide to GAISSF™	Board-level oversight questions	COM-018 helps the CISO build evidence-backed management answers; it does not replace the board guide or turn board

Publication	Role retained in COM-018	Consistency rule
		questions into assessment criteria.
COM-011 — When AI Breaks: The ODA3 Guide to AI Incident Response	Short incident-response orientation	COM-018 designs readiness and executive interfaces; active response remains outside this guide.
COM-012 — Classifying AI Incidents: A Practical Guide to UAIF™	UAIF™ classification / uncertainty discipline	Local observation or triage does not create a UAIF `Observation` record type; technical routing does not create a legal conclusion.
COM-013 — AI Incident Response Field Guide: Using UAIF™ and AI-IRF™ for Real-World AI Incidents	Full classification/response integration	Facts, hypotheses, provider representations, unknowns and specifically checked-but-not-observed conditions remain distinguishable; response may evolve while classification remains incomplete.
COM-014 — Securing Physical AI: A Practitioner's Guide to PAI-SF™	Physical-AI practitioner boundary	D9 continuity, PAI-SF™ supersession/expansion, meaningful human authority and safety/return-to-service separation govern COM-018 physical-AI language.
COM-015 — Building an AI Assurance Evidence Pack	Evidence preparation	COM-018 may consume evidence state, but COM-015 owns evidence indexing, provenance, evidence-source states, limitations, contradictions and handoff preparation. The public practitioner journey ends at this evidence-preparation handoff; later assessment activity is separately governed and authorized.
COM-017 — Implementing GAISSF™	Detailed control implementation	COM-018 sets CISO intent, ownership, challenge and portfolio expectations; COM-017 owns decomposition, mechanism implementation and implementation-side verification.

Where older practitioner shorthand differs from a newer controlling/current clarification, the current controlling source or current FINAL practitioner boundary governs.

Category and Management Models

Category: Operational AI Assurance.
CISO operating lens: AI Security Assurance.

Informative CISO Assurance Loop:
Know → Bound → Control → Evidence → Observe → Respond → Challenge → Report → Improve

Three management planes:

- **System Assurance Plane** — system boundary, controls, implementation, evidence, observability and incident readiness.
- **Portfolio Assurance Plane** — enterprise inventory, providers, common dependencies, exceptions, remediation, material change and concentration exposure.
- **Executive Assurance Plane** — risk decisions, challenge, reporting, uncertainty, claims and escalation.

These are management aids, not GAISSF domains, assessment categories, maturity levels or certification tiers.

Executive Navigation — Use COM-018 Without Reading It Linearly

COM-018 is designed as a reference guide, not as a mandatory page-by-page programme.

The 10-Minute CISO Orientation

If time is limited, read these sections first:

1. **Section 1 — Executive Summary**
2. **Section 5 — The CISO Assurance Gap**
3. **Section 6 — What the CISO Owns — and What Is Shared**
4. **Section 14 — Build the AI Assurance Portfolio**
5. **Section 45 — Assurance Is Demonstrated, Not Asserted**
6. **Section 56 — Design Incident Readiness Before the Incident**
7. **Section 76 — Design the CISO AI Assurance Dashboard**
8. **Section 83 — Report to the Board and Risk Committee**
9. **Section 95 — Continue, Constrain, Degrade, Suspend or Retire**
10. **Sections 96–98 — CISO practitioner tools**

This route gives the management spine. Move into the specialist sections only where the portfolio or current decision requires it.

Role-Based Routes

Role / responsibility	Start here	Then use
CISO / Deputy CISO	1–8	9–18, 67–86, 94–98
AI Security Lead	19–30	31–44, 56–74, 87–93
Security Architect	21–30	31–44, then COM-017
SOC / CSIRT leader	56–66	92, then COM-011 / 012 / 013
GRC / enterprise risk	9–18	45–55, 67–86, 94, 97–98

Role / responsibility	Start here	Then use
Supplier / third-party security	22, 26, 28	38, 49, 61, 69, 81, 88
Engineering / platform leadership	19–30	31–44, then COM-017
Internal assurance / audit interface	45–55	71, 83–86, 94, then COM-015; any later assessment activity follows the separately authorized route applicable to the engagement
Safety / OT / physical-AI interface	29, 42	63, 93, then COM-014 / PAI-SF™
Board / risk-committee preparation	14, 45, 75–86	94, 98

Situation-Based Routes

Situation	Recommended COM-018 route
We are starting an enterprise AI security-assurance programme	9–18 → 19–30 → 31–44 → 67–86 → 96–98
We cannot confidently identify our AI estate	19–30 → 77 → 97
A new material AI system is proposed	19–31 → 87 → 44
We are onboarding a hosted model / AI provider	22, 26, 28, 38, 49, 61, 69, 81 → 88
We are deploying an agent or tool-using AI system	24, 25, 27, 34, 40, 56, 58, 60, 62 → 89
We are deploying RAG over enterprise knowledge	23, 27, 37, 40, 58 → 90
A model, tool, provider or permission changed materially	30, 52, 68 → 91
We suspect an AI security incident	56–66 → 92 → COM-011 / 012 / 013
The AI system can affect the physical world	29, 42, 63 → 93 → COM-014 / PAI-SF™
The board asks whether our AI is secure	14, 45, 75–86 → 94
Audit / customer / regulator asks for assurance	45–55, 83–86 → 94 → COM-015 → separately authorized assessment route, where applicable
The current security basis is no longer trusted	30, 60–65, 68, 72 → 95

Decision-Playbook Map

Sections **87–95** are the fastest path when the immediate question is a management decision rather than programme design:

- **87** — new AI system or major use case;
- **88** — third-party AI provider;
- **89** — agentic AI;
- **90** — RAG / enterprise knowledge;

- **91** — material model, tool or provider change;
- **92** — suspected AI security incident;
- **93** — high-consequence / physical AI;
- **94** — assurance request from board, audit, customer or regulator;
- **95** — continue, constrain, degrade, suspend or retire.

These are decision aids, not mandatory ODA3 gates or assessment states.

Practitioner Tool Map

- **Section 96 — CISO AI Security Assurance Charter** — use when authority, interfaces or escalation are unclear.
- **Section 97 — AI Assurance Portfolio Register** — use when the CISO needs enterprise visibility across systems, providers, evidence, exceptions, incidents and material change.
- **Section 98 — CISO Decision Rights and Claims Checklist** — use immediately before a consequential security decision or an assurance statement to executives, auditors, customers or external parties.

Part I — The CISO Mandate for AI Security Assurance

Establish what AI security assurance means for the CISO, how it relates to Operational AI Assurance, and where CISO accountability begins and ends.

1. Executive Summary

AI security assurance is the CISO's ability to move beyond the statement that an AI system is "secure" and show what is actually known about the system, which controls are expected, who owns them, what evidence supports their operation, which dependencies and uncertainties remain, and how the organization will respond when confidence is lost. The objective is not perfect certainty. It is a bounded, evidence-backed management position that can survive technical challenge, executive scrutiny and material change.

CISO operating questions / actions:

- Know which AI systems and AI-enabled business processes materially affect the security posture.
- Make ownership, provider responsibility, delegated authority and high-consequence paths visible.
- Require controls to produce usable operational evidence rather than only design documentation.
- Make incident reconstruction and containment part of system design, not an emergency improvisation.
- Report material unknowns, exceptions and evidence limitations alongside positive indicators.

COM-018 therefore treats the CISO as the executive security integrator of the assurance system—not the sole owner of every AI risk. Business owners, engineering, governance, privacy/legal, internal assurance, providers and safety functions retain their own decision authorities.

---

2. Purpose and Audience

This guide is written for security leaders who need to turn AI security from a collection of projects into an enterprise assurance capability. It assumes the reader already has conventional cybersecurity functions—identity, cloud, application security, SOC/CSIRT, supplier security, GRC, resilience and change management—and needs to determine where AI changes the operating model.

CISO operating questions / actions:

- Use the guide to establish CISO decision rights and escalation routes.
- Use it to build a portfolio-level view across AI systems rather than managing assurance one assessment at a time.
- Use it to determine what evidence and incident-readiness capabilities security leadership should expect.
- Use the detailed WHP guides when the task becomes control implementation, evidence-pack preparation, active incident response or physical-AI engineering; any later assessment activity follows the separately governed and authorized route applicable to the engagement.

Primary readers are CISOs, Deputy CISOs, AI Security Leads and Security Architecture leaders. Adjacent functions should use the guide to understand where they interface with CISO assurance rather than to transfer their own accountability to security.

---

3. AI Security Assurance in Plain Terms

For COM-018, AI security assurance means justified confidence—supported by bounded evidence—that AI-related security controls and response capabilities are appropriately designed, implemented, operating and governed for the stated system and period. The word “assurance” is important because it requires the organization to distinguish what it believes from what it can demonstrate.

CISO operating questions / actions:

- State the claim narrowly enough that evidence can actually support it.
- Identify whether the support comes from enterprise observation, provider representation, testing, operating records or professional judgement.
- Expose conditions and limitations that would invalidate or narrow the claim.
- Preserve unknowns rather than forcing a positive/negative conclusion where the information is incomplete.

AI security assurance is not a guarantee that harm will not occur. In this guide, the CISO’s assurance position is a management position, not an independent audit/assurance opinion or controlled GAISSF assessment determination. It is also not automatically a certification, legal-compliance conclusion, safety approval or prediction of future system behavior.

---

4. Relationship to Operational AI Assurance

Operational AI Assurance remains the broader ODA3 category: the discipline of determining whether AI systems are governed, secure, evidenced, monitored, incident-ready and

recoverable in the environments where they operate. COM-018 uses AI Security Assurance as the CISO's operating lens inside that broader discipline.

CISO operating questions / actions:

- Use Operational AI Assurance when describing the enterprise category and the full governance-to-evidence-to-response problem.
- Use AI Security Assurance when discussing the security leader's responsibilities, security controls, evidence expectations and challenge role.
- Do not present the two terms as competing categories.
- Keep non-security ownership—privacy, legal, business risk, model risk, safety and sector obligations—visible even when security leads the assurance coordination.

This hierarchy lets the CISO operate within a coherent enterprise assurance architecture without claiming that cybersecurity owns every dimension of AI assurance.

---

5. The CISO Assurance Gap

The assurance gap appears when management can point to an AI policy, a vendor review, a red-team report or a security product but cannot explain what currently operates for the production system and what evidence supports that position. AI enlarges this gap because behavior can depend on changing models, prompts, retrieval sources, providers, tools, delegated authority and downstream actions.

CISO operating questions / actions:

- Ask which system/version/environment a security statement actually covers.
- Ask whether the control is designed, implemented, operating and observable—or merely documented.
- Ask how provider changes or opaque services could invalidate the assurance position.
- Ask whether incident responders could reconstruct consequential AI actions today.
- Ask what material unknowns are being hidden by a simple green dashboard indicator.

The CISO's job is not to eliminate uncertainty. It is to prevent uncertainty from being silently converted into confidence.

---

6. What the CISO Owns — and What Is Shared

CISO accountability is strongest around security architecture, security-control expectations, cyber incident readiness, security exceptions, security evidence and independent challenge. It is shared with other authorities whenever the decision also affects business risk, privacy, legal obligations, model governance, financial risk, human safety or sector approval.

CISO operating questions / actions:

- Own the security-assurance operating model and security challenge function.
- Require material AI systems to have identifiable security owners and interfaces with business/system owners.

- Escalate risks that exceed security authority rather than accepting them by default.
- Keep safety, legal, privacy and regulatory decision rights outside the CISO function unless the organization has explicitly assigned them there.

A useful test is: the CISO should be able to stop pretending that a risk is resolved, but may not be the person authorized to accept the underlying business, legal or safety risk.

---

7. GAISSF Ecosystem Position

COM-018 uses the GAISSF Ecosystem as an integrated set of sibling frameworks, selected by objective and scope rather than arranged as a maturity sequence.

CISO operating questions / actions:

- Use GAISSF™ for the AI safety/security control and assurance baseline.
- Use UAIF™ when a potentially material AI event needs structured classification and evidence discipline.
- Use AI-IRF™ for structured AI incident preparation and response.
- Use PAI-SF™ where AI-enabled sensing, inference, autonomy or command can create or materially influence physical consequence.

More than one framework can apply to the same system. COM-018 does not create a new CISO-specific GAISSF profile or modify the authoritative framework architecture.

Framework maturity and technical-publication state remain separate from the publication status of this guide. UAIF™ and AI-IRF™ are **Final Publication v1.0**; PAI-SF™ retains its current 12-domain / 41-control public architecture. Publication of COM-018 does not change those underlying states.

---

8. How to Use This Guide

COM-018 is a reference and decision guide. A CISO should not need to read all 100 sections before acting. Start with the immediate management problem, use the relevant Part, and move to the specialist publication once the work drops below the CISO decision layer.

CISO operating questions / actions:

- Mandate/ownership problem → Part II.
- Unknown AI estate or boundary → Part III.
- Control-architecture question → Part IV, then COM-017 for implementation depth.
- Evidence problem → Part V, then COM-015.
- Assessment question → determine the separately authorized assessment route applicable to the engagement; COM-018 does not define or activate that route.
- Active incident → COM-011/012/013 and the applicable UAIF™ / AI-IRF™ process.
- Physical-AI depth → COM-014 / PAI-SF™.

The loop used throughout the guide—Know → Bound → Control → Evidence → Observe → Respond → Challenge → Report → Improve—is informative and recursive. Material change can send the organization back to any earlier point.

For practical use, choose the route that matches the current decision:

- **Need enterprise visibility?** Start with Parts II–III.
- **Need to understand the security architecture expectation?** Start with Part IV.
- **Need to defend a management claim?** Start with Part V.
- **Need incident readiness?** Start with Part VI.
- **Need operating governance or remediation oversight?** Start with Part VII.
- **Need board reporting or metrics?** Start with Part VIII.
- **Need to make a specific CISO decision now?** Start with Part IX.
- **Need a reusable management template?** Start with Part X.

No route changes framework applicability, assessment state or authorization.

Part II — Establish the Enterprise AI Assurance Operating Model

Use this Part when: the organization has AI activity but unclear security authority, ownership, escalation or portfolio governance.

Working output: a CISO mandate/charter, decision-rights map, ownership model, portfolio structure and exception route.

Part checkpoint: the CISO should be able to identify who owns the system, who owns material controls, who can accept unresolved risk, and where security authority stops.

Turn AI security assurance into an enterprise operating capability with mandate, decision rights, ownership, portfolio governance, exceptions and resourcing.

9. Establish the CISO AI Assurance Mandate

AI assurance fails when the CISO is expected to answer enterprise questions without authority to obtain system information, challenge unsafe design or escalate unresolved evidence gaps. The mandate should therefore establish access and decision interfaces before metrics or tooling are selected.

CISO operating questions / actions:

- Define which AI systems and business processes can enter the security-assurance portfolio.
- Give security authority to request architecture, provider, identity, logging and incident-readiness information.
- Define when security may require remediation, constrain exposure or escalate to a higher risk authority.
- State where business, legal, privacy, model-risk and safety decisions remain outside security authority.

A mandate is not a declaration that the CISO owns all AI risk. It is the basis for effective security challenge and coordination.

---

10. Create an AI Security Assurance Charter

The charter converts the mandate into an operating agreement. It should be short enough to be used and precise enough to prevent security, governance and engineering teams from assuming different definitions of assurance.

CISO operating questions / actions:

- Purpose and scope of the CISO assurance programme.
- In-scope system and portfolio criteria, with a route for scope changes.
- Decision rights, escalation routes and risk-acceptance interfaces.
- Interfaces to define with system owners, AI governance, engineering, providers, SOC/CSIRT, GRC, privacy/legal, internal assurance and safety functions.
- Reporting audiences, review triggers and claims boundaries.

Section 96 provides a public-safe charter template. The charter is an internal management artifact, not an ODA3 assessment or readiness instrument.

---

11. Map Governance Interfaces

AI assurance is a cross-functional system. The CISO should deliberately map interfaces rather than relying on job titles to imply ownership.

CISO operating questions / actions:

- CIO/CTO: platforms, architecture, engineering and operational ownership.
- AI governance/model risk: use-case governance, model lifecycle and non-security AI risk.
- Legal/privacy: legal obligations, data rights, privilege and reporting decisions.
- Procurement/supplier risk: contractual controls, provider evidence and exit provisions.
- SOC/CSIRT: monitoring, triage, evidence preservation, containment and recovery.
- Internal assurance/audit: independent challenge and use of evidence.
- Safety/OT: physical consequence, safe-state and return-to-service authority where applicable.

The interface map should show who decides, who supplies evidence, who operates and who challenges. A RACI alone is insufficient if decision authority remains ambiguous.

---

12. Define Decision Rights

A mature programme distinguishes the right to advise from the right to decide. CISO decision rights should be explicit for normal operation, exceptions, incidents and major change.

CISO operating questions / actions:

- Who can approve a security design or security exception?
- Who can block or constrain deployment for an unresolved security condition?

- Who may revoke agent/tool authority or privileged credentials?
- Who may initiate cyber containment during an incident?
- Who accepts residual business risk when security recommends against operation?
- Who makes safety/operational return-to-service decisions for physical AI?

Where both authorities govern the decision—for example security closure and safety/operational return-to-service—record both. Do not collapse them into a single "security approval."

---

13. Assign System and Control Ownership

Assurance becomes fragile when a system is 'owned' by a committee but no named role is accountable for the actual service or control. Each material AI system should have an identifiable system owner and a traceable set of control owners.

CISO operating questions / actions:

- System owner: accountable for business use, scope and operation.
- Control owner: accountable for a particular security/control capability.
- Operator/custodian: runs the mechanism or process.
- Platform/provider owner: manages a shared dependency or third-party relationship.
- CISO/security challenge: reviews whether the representation is supportable and whether material gaps are escalated.

Shared ownership is valid, but a useful shared-responsibility record should still identify each party's mechanism, evidence visibility and failure responsibilities.

---

14. Build the AI Assurance Portfolio

The portfolio is the CISO's enterprise-level view of AI security assurance. It should connect systems to owners, providers, consequence paths, material controls, evidence state, incidents, exceptions and change exposure without reducing them to a single score.

Any portfolio fields such as `planned`, `in progress`, `implemented`, `blocked`, `exception`, `unknown` or similar organization-defined labels remain local management states. They are not GAISSF control determinations, assessor dispositions, readiness categories or certification states.

CISO operating questions / actions:

- Group systems by consequence and operating context rather than only by business unit.
- Highlight unowned systems, opaque providers and common dependencies.
- Track material implementation/evidence gaps and unresolved exceptions.
- Link incidents and material changes back to the affected systems and controls.
- Maintain a visible set of high-consequence or physical-AI systems requiring additional governance.

The portfolio answers 'where should CISO attention go next?' It is not a substitute for formal system-level applicability analysis or assessment results.

15. Set Exception and Risk-Acceptance Routes

Exceptions are part of real security operations. The assurance problem is not that exceptions exist; it is that temporary gaps become invisible, ownerless or repeatedly renewed without renewed analysis.

CISO operating questions / actions:

- Record the affected system/control and why the normal mechanism cannot currently operate.
- Document the compensating measure and its limitations.
- Identify monitoring needed while the exception exists.
- Set a review/change trigger appropriate to the context rather than a universal COM-018 deadline.
- Route residual risk to the person or body authorized to accept it.

An exception record is not evidence that the underlying control is satisfied. It is evidence that a known gap is being governed.

---

16. Use a Federated Operating Model

Large organizations should not centralize every AI security decision in one specialist team. A federated model lets a central CISO function define assurance expectations while product, platform and business teams implement them within controlled boundaries.

CISO operating questions / actions:

- Central team: principles, common control interpretation, portfolio visibility, challenge and reporting.
- Business/product teams: system-specific implementation and operating evidence.
- Platforms: reusable identity, logging, gateway, policy, model-management and provider controls.
- Regional functions: local obligations, operating context and escalation.
- Internal assurance: independent review separate from implementation ownership.

Federation only works if the central team can see material deviations and the local teams know which decisions they may make without escalation.

---

17. Reuse Existing Cybersecurity and GRC Structures

AI security assurance should attach to established enterprise control planes wherever they remain effective. Reuse reduces duplication and increases the chance that controls will operate continuously rather than only during AI reviews.

CISO operating questions / actions:

- Extend existing IAM and privileged-access processes to AI services, agents and tools.
- Use established SDLC/DevSecOps and change processes for AI application/model changes where appropriate.
- Extend supplier risk to models, AI APIs, data sources, tools and hosted inference.
- Feed AI-relevant telemetry into existing SOC/CSIRT workflows where it remains decision-useful.
- Use GRC to manage ownership, exceptions and obligations without assuming the GRC platform itself proves operating effectiveness.

The CISO should prefer extension over duplication, but should not force AI-specific risk into a legacy control that cannot represent it.

---

18. Resource, Skills and Capacity Planning

The CISO needs enough capability to challenge AI systems even when implementation remains distributed. Capability planning should therefore cover expertise and decision capacity, not only headcount.

CISO operating questions / actions:

- AI architecture and model/application security.
- Agentic AI, delegated authority and tool/protocol security.
- AI supply-chain/provider security and evidence negotiation.
- Evidence engineering, observability and incident reconstruction.
- AI incident readiness and response integration.
- Physical-AI coordination with OT/safety specialists where applicable.

Organizations may build, buy or partner for capability, but third-party delivery remains subject to applicable licensing and authorization boundaries. External expertise does not transfer accountability away from the enterprise.

Part III — Know the AI Estate and Security Boundary

Use this Part when: the enterprise cannot confidently answer what AI exists, which providers/models/tools are involved, or what the real security boundary is.

Working output: an enterprise AI-system inventory and bounded dependency/authority view suitable for security management.

Part checkpoint: material AI systems should be identifiable enough to map owners, providers, identities, data/retrieval, tools/actions, dependencies, change and physical consequence.

Create enough visibility into AI systems, providers, data, agents, identities, dependencies, interfaces and physical consequence paths to support meaningful security decisions.

19. Build the AI System Inventory

A CISO cannot assure what the organization cannot identify. The inventory should be broad enough to capture material AI use, but it should not pretend to be the formal GAISSF applicability decision itself.

CISO operating questions / actions:

- Record the business process/use case, owner and operating environment.
- Identify models/services/providers and deployment route.
- Capture data/retrieval sources, tools/actions and key integrations.
- Flag high-consequence, regulated or physical-effect characteristics.
- Record lifecycle state and last material change.

Inventory completeness is a management objective, not proof that every use case has been found. Unknown/shadow AI should remain a visible residual uncertainty.

---

20. Find Shadow, Embedded and Indirect AI

Traditional asset discovery often misses AI that arrives as a feature of SaaS, a developer assistant, an API option, a security product capability or an automated workflow. These systems can create material data, identity or decision paths even when no team says it ‘deployed a model.’

CISO operating questions / actions:

- Review major SaaS and cloud platforms for AI features and changes.
- Inspect developer tooling, IDE assistants, code agents and automation platforms.
- Look for external model/API calls in repositories, gateways and network telemetry where lawful and feasible.
- Review procurement and expense channels for AI services.
- Use business-process discovery for AI embedded in operations rather than relying only on technical scans.

Discovery methods should respect privacy, employment and monitoring constraints. The objective is bounded visibility, not indiscriminate surveillance.

---

21. Define the Working System Boundary

Security decisions require a boundary that reflects the real system rather than only the model. The boundary should include the components and actors that can affect the security claim or the consequence of failure.

CISO operating questions / actions:

- Business process and users.
- Application/orchestration layer.
- Models/providers/routing.
- Data, retrieval and knowledge sources.

- Agents, tools, connectors and downstream systems.
- Identity/credential infrastructure.
- Monitoring, logging and incident interfaces.
- Physical-effect path where relevant.

The working boundary is a CISO management view. Formal applicability and assessment scope remain governed by the controlling framework/process.

---

22. Track Models, Providers and Versions

A model or provider change can alter behavior, data handling, available tools, safety controls or evidence access without a conventional software deployment. Version visibility is therefore part of the assurance boundary.

CISO operating questions / actions:

- Record model/service identity and hosting/provider context.
- Track routing or fallback that can switch the effective model.
- Preserve version/configuration information needed for reconstruction.
- Require provider-change visibility where contractually and technically feasible.
- Treat unannounced or opaque changes as an assurance limitation, not automatically as proof of failure.

The CISO does not need every model parameter. The requirement is enough identity and change visibility to manage security assumptions and incidents.

---

23. Map Data, Retrieval and Knowledge Sources

AI security depends on what data the system can access and what context can influence model behavior. RAG and connector-based systems turn data access into a runtime control surface.

CISO operating questions / actions:

- Identify major source classes and access rules.
- Map ingestion/indexing and source-approval paths.
- Understand retrieval authorization and tenant/user separation.
- Track connector permissions and delegated access.
- Identify source freshness, poisoning/integrity and instruction-content risks where relevant.

Do not assume that citations or source labels alone prove trustworthy retrieval. The security question is whether the source, access path and context assembly are governed.

---

24. Map Agents, Tools and External Actions

Agentic systems should be mapped as authority-bearing systems, not only as model prompts. The key assurance question is what the system can cause to happen.

CISO operating questions / actions:

- Agent/service identity and sponsoring principal.
- Objectives and orchestration logic.
- Tools/connectors and permitted actions.
- Credentials and delegated authority.
- Approval points, limits and destinations.
- Memory/state and persistence.
- Transactions or externally consequential actions.
- Monitoring, revocation and rollback surfaces.

An agent with read-only research tools and an agent able to change production infrastructure are not the same assurance problem even if they use the same model.

---

25. Map Identities, Credentials and Privilege

AI introduces machine identities, delegated credentials and service-to-service authority that may bypass assumptions made for human users. The CISO should connect these paths to existing identity governance.

CISO operating questions / actions:

- Use identifiable service/agent identities where architecture permits.
- Minimize long-lived secrets and broad inherited credentials.
- Separate human approval from credential delegation where appropriate.
- Review tool/API scopes, destinations and transaction authority.
- Ensure rapid revocation is possible for high-consequence agents or integrations.

Least privilege remains relevant, but AI often requires an additional question: what is the least agency the system needs to perform the approved task?

---

26. Map AI and Software Supply-Chain Dependencies

AI systems inherit risks from conventional software plus models, registries, tool servers, plugins, datasets and hosted services. Dependency mapping should therefore cross the model/software boundary.

CISO operating questions / actions:

- Models and model hubs/registries.
- Libraries, packages and runtime components.
- External model/inference APIs.
- MCP/tool servers, plugins and connectors.
- Datasets, embeddings and retrieval services.

- Cloud/compute/platform dependencies.
- Update and artifact-delivery channels.

The portfolio should highlight common dependencies whose failure or compromise could affect many AI systems at once.

---

27. Map Interfaces, Protocols and Trust Boundaries

The most consequential AI security failures often occur at interfaces—where instructions, data, tool definitions, credentials or state cross a boundary that teams implicitly trust.

CISO operating questions / actions:

- User ↔ application/model.
- Application ↔ provider/model API.
- Agent ↔ tool/MCP server.
- Retrieval service ↔ knowledge source.
- Model output ↔ downstream automation.
- Provider ↔ enterprise telemetry/support.
- Digital control ↔ physical actuation where applicable.

A boundary map should identify what is trusted, how that trust is established, and what becomes visible when the assumption fails.

---

28. Make Provider Opacity and Shared Responsibility Explicit

Provider dependence is not automatically a control weakness, but opaque responsibility is. The CISO should distinguish who operates the mechanism from who can observe or evidence it.

CISO operating questions / actions:

- Enterprise-controlled: mechanism and evidence primarily under enterprise control.
- Provider-controlled: provider operates the mechanism; enterprise evidence may be limited.
- Shared: both parties carry material responsibilities.
- Visibility-limited: the enterprise cannot directly observe a material element.

Keep this **implementation-responsibility** view separate from COM-015's **evidence-source** states. For evidence preparation, the current public states are **Customer-observed / Provider-provided / Shared-interface / Provider-only-inaccessible / Unknown**. These evidence-source states describe where evidentiary support comes from; they do not replace the implementation-responsibility descriptors above and are not control ratings or assessment outcomes.

A provider statement can be useful evidence but remains a provider representation unless independently observed or otherwise corroborated.

---

29. Identify Physical-Effect Paths

When AI can influence movement, force, process state, infrastructure state or another physical condition, the CISO security boundary should extend beyond the software stack. Security containment can itself create physical risk.

CISO operating questions / actions:

- Trace the operational consequence chain far enough to understand security authority and physical effect. A useful detailed sequence from the current physical-AI practitioner set is: **environment → sensing/perception → inference → autonomy state → command authorization → command → actuator state → safety-controller state → operator action → physical outcome → recovery/revalidation**.
- Identify perception/sensing and environmental assumptions.
- Map autonomy/authority boundaries and command authorization.
- Map command, actuator and safety-controller states where relevant.
- Identify override, safe/degraded modes, operator intervention and local fallback.
- Identify the safety/operational authority required for containment, recovery and revalidation.

GAISSF D9 remains the continuity point; detailed treatment hands off to PAI-SF™. COM-018 does not replace functional safety, sector approval or return-to-service authority.

---

30. Track Material Change

The assurance position is time-bound. A material change can invalidate assumptions even when the formal system name is unchanged.

CISO operating questions / actions:

- Model/provider/version or routing change.
- Prompt/orchestration/tool change.
- New data or retrieval source.
- Permission/credential/identity change.
- Infrastructure or logging change.
- Business use, user population or jurisdiction change.
- Autonomy or physical-operating-environment change.
- Incident, near miss or newly discovered dependency.

The CISO should define change triggers that reopen the right layer of analysis. Not every change requires a full reassessment; materiality and the controlling process determine the response.

Part IV — Build the Enterprise AI Security Control Architecture

Use this Part when: the CISO needs to translate the applicable GAISSF baseline into enterprise security expectations without becoming the implementation team.

Working output: a CISO-level control architecture and a clear handoff into COM-017 for implementation detail.

Part checkpoint: existing enterprise controls are reused where adequate, AI-specific extensions are visible, and ownership/evidence/incident-readiness expectations are explicit.

Connect the applicable GAISSF control baseline to existing enterprise security architecture and identify the AI-specific extensions that require CISO attention.

31. Establish the Applicable Security-Control Baseline

COM-018 should not invent a separate CISO control catalogue. The enterprise should use the authoritative GAISSF scope/applicability process to determine which requirements apply to the system and then translate the resulting control baseline into security architecture and ownership.

CISO operating questions / actions:

- Confirm the governed scope and applicable control baseline.
- Identify security-owned, shared and non-security-owned obligations.
- Record where applicability is uncertain or dependent on unresolved scope facts.
- Keep sector/legal obligations as additive overlays rather than assuming GAISSF automatically satisfies them.

The CISO assurance portfolio may summarize applicable-control state, but it does not replace the controlling applicability record.

---

32. Reuse Existing Security Controls First

A strong CISO programme begins by asking which existing enterprise controls already carry the intent. Reuse reduces fragmentation and produces more credible evidence because established controls already operate at scale.

CISO operating questions / actions:

- IAM/privileged access.
- Network/cloud workload security.
- Secure SDLC/DevSecOps.
- Data access/governance.
- Logging/SIEM/SOC.
- Supplier/security review.
- Resilience/change management.

Reuse should be evidence-based. A similarly named legacy control is not automatically equivalent if it does not cover AI-specific authority, behavior, provider or evidence conditions.

---

33. Identify the AI-Specific Extension

The CISO should concentrate scarce expertise on the points where ordinary controls stop being sufficient. These extensions are usually caused by the behavior or operating context of the AI system rather than by the existence of an AI label.

CISO operating questions / actions:

- Model/provider version and behavioral change.
- Prompt/instruction and retrieval manipulation.
- Agent/tool authority and delegated credentials.
- Provider opacity and mutable external services.
- AI-specific incident reconstruction.
- Human oversight effectiveness.
- Physical consequence paths.

The extension may be a new control, a stronger implementation of an existing control, better evidence, or a new decision interface.

---

34. Identity, Access and Least Agency

AI security architecture should constrain not only who can access a system but what an AI-enabled component can decide or cause. Least agency extends least privilege into delegated action.

CISO operating questions / actions:

- Dedicated identities for material agents/services where feasible.
- Tool allowlists and scoped API permissions.
- Parameter, destination, transaction-value or time constraints where relevant.
- Context-specific approvals for high-consequence actions.
- Credential isolation and short-lived delegation where feasible.
- Rapid revocation/circuit-breaker paths.

Human approval is not universally required for every autonomous action. The level and placement of approval should reflect consequence, reversibility, operational context and existing control design.

---

35. Runtime Security and Adversarial Defense

Runtime controls should address the fact that valid users and approved models can still produce unsafe or unauthorized effects when instructions, context, tools or outputs are manipulated.

CISO operating questions / actions:

- Input/instruction abuse and prompt-injection pathways.
- Tool misuse and authority escalation.
- Anomalous sequences or actions where detectable.

- Unsafe routing or fallback.
- Output-to-action gating where consequence requires it.
- Runtime policy enforcement and emergency revocation.

The CISO should avoid claiming that a single gateway, filter or guardrail provides universal AI security. Runtime defense is layered and system-specific.

---

36. Secure AI Development and MLOps

AI application security should be integrated into existing engineering governance rather than delegated to periodic model tests.

CISO operating questions / actions:

- Approved source repositories and change control.
- Model/artifact provenance and integrity where relevant.
- Environment and secret separation.
- Evaluation/security testing tied to material change.
- CI/CD and deployment authorization.
- Rollback and last-known-good restoration.
- Traceability between deployed version and reviewed artifacts.

Model evaluation is one input to assurance. It does not replace secure engineering, access control, supplier assurance, operational evidence or incident readiness.

---

37. Secure RAG and Enterprise Knowledge Paths

RAG can expand the effective attack surface from a model endpoint to every source, connector and authorization path that can influence context.

CISO operating questions / actions:

- Approve and identify knowledge sources.
- Protect ingestion/indexing integrity.
- Apply authorization before retrieval where architecture permits.
- Limit connector/data access to the task and user context.
- Separate untrusted instructions from authoritative context where feasible.
- Track source/version changes and re-indexing.

The CISO should also ask whether incident responders can reconstruct which source/context materially influenced a consequential output.

---

38. Secure Providers and the AI Supply Chain

Provider assurance should move beyond the one-time security questionnaire. The CISO needs an operating view of provider identity, service boundaries, change, incident support and evidence access.

CISO operating questions / actions:

- Approved services/endpoints and data-handling terms.
- Provider change-notification and version transparency where available.
- Subprocessors/dependencies relevant to the security claim.
- Incident escalation and evidence request channels.
- Business continuity, fallback and exit options.
- Known visibility limitations recorded in the assurance position.

Provider attestations can support assurance but should not be presented as enterprise-observed operating evidence when the enterprise cannot see the underlying mechanism.

---

39. Protect Output, Human Review and Downstream Use

AI output becomes more consequential when another person, workflow or system acts on it. The CISO should therefore examine the downstream authority path, not only the safety quality of the output.

CISO operating questions / actions:

- Who or what consumes the output?
- What decision/action can follow automatically?
- What approval or verification exists for high-consequence use?
- Can the human reviewer understand, challenge and stop the action?
- Does the system preserve enough context for later reconstruction?

Nominal ‘human in the loop’ language is weak assurance unless meaningful human authority is credible across the current practitioner dimensions: **Awareness / Comprehension / Time / Authority / Effectiveness**. A review step shown on a diagram does not prove that the human could understand, challenge, stop or change the consequential action in practice.

---

40. Design Observability for Assurance

Observability should answer security questions, not simply maximize telemetry. The CISO should know which signals are needed to detect change, investigate incidents and support control evidence.

CISO operating questions / actions:

- System/model/provider/version identity.
- User/agent identity and credential context.
- Tool/connector invocation and consequential actions.
- Configuration/policy changes.

- Relevant retrieval/source identifiers.
- Containment/revocation and recovery events.
- Provider incident/change notifications.

COM-018 does not require universal logging of all prompts or outputs. Collection should be proportionate and constrained by privacy, privilege, security and operational needs.

---

41. Design Resilience, Fallback and Rollback

A secure AI system needs a controlled response when the primary model, provider, tool or evidence path cannot be trusted. Resilience is therefore about preserving acceptable control, not merely keeping the service available.

CISO operating questions / actions:

- Define last-known-good versions/configurations where feasible.
- Design rollback or provider/model fallback paths.
- Consider data/schema and transaction compatibility.
- Define degraded modes that reduce authority or functionality.
- Know when suspension is safer than continued operation.
- Validate recovery before restoring normal authority.

For physical AI, a digital shutdown or network isolation action may itself be unsafe; specialist safety/operational authority remains required.

---

42. Physical-AI Security Handoff

The CISO should treat physical AI as an extension of the security boundary into sensing, autonomy, command and effect. The detailed control treatment belongs in PAI-SF™, not in a generic enterprise AI-security checklist.

CISO operating questions / actions:

- Trigger the handoff when AI can cause, influence or fail to prevent material physical action.
- Preserve GAISSF D9 continuity.
- Bring OT/safety/engineering owners into the assurance model.
- Ensure cyber containment decisions account for safe/degraded behavior.
- Keep functional safety, product safety and sector approval distinct from security assurance.

COM-018 can require the interface to exist; it does not authorize a physical system to return to service.

---

43. Manage Exceptions and Compensating Controls

A CISO assurance programme should make exceptions visible at both system and portfolio level. The aim is to prevent temporary security compromises from becoming the organization's undocumented steady state.

CISO operating questions / actions:

- Affected system/control and scope.
- Reason for the exception.
- Compensating mechanism and residual limitation.
- Monitoring/change triggers.
- Owner and risk authority.
- Remediation/dependency plan.

A compensating control should be challenged for actual operating coverage; its existence in an exception record does not prove equivalent control effectiveness.

---

44. Handoff to COM-017 — Implementing GAISSF™

COM-018 should stop before it becomes an implementation handbook. Once the CISO has established the applicable control baseline, ownership, enterprise architecture expectations and management constraints, detailed implementation belongs in COM-017.

CISO operating questions / actions:

- COM-018: what the CISO requires, funds, challenges, monitors and escalates.
- COM-017: how practitioners decompose systems, translate requirements, select mechanisms, integrate controls, verify implementation, operate and change them.
- COM-015: how resulting artifacts are prepared as evidence. The public practitioner journey ends at this evidence-preparation handoff; later assessment activity is separately governed and authorized.

This separation prevents CISO guidance from accidentally becoming a parallel GAISSF implementation standard.

Part V — Make Security Assurance Evidentiary

Use this Part when: management claims are stronger than the evidence basis, provider evidence is opaque, or the organization is preparing for audit/assessment/executive scrutiny.

Working output: an evidence-expectation model that preserves scope, period, provenance, provider source, contradictions and limitations.

Part checkpoint: the CISO can distinguish implementation records from defensible evidence and explain what remains unknown without inventing an assessment conclusion.

Ensure that CISO confidence is traceable to bounded evidence, provenance, limitations, contradictions and truthful claims rather than policy or tooling assertions alone.

45. Assurance Is Demonstrated, Not Asserted

The CISO should be able to trace a material security statement to the evidence that supports it. A policy, tool purchase or responsible owner's statement can be relevant, but each has a different evidentiary meaning.

CISO operating questions / actions:

- Ask what claim is being made.
- Identify which system, environment, version and period the claim covers.
- Identify the evidence source and its independence.
- Record material limitations, contradictions and unknowns.
- Separate evidence of design from evidence of operation.

The maxim for the programme is simple: assurance is demonstrated, not asserted.

---

46. Build Evidence into the Security Architecture

Evidence is most credible when it is generated naturally by controlled operation. A late collection exercise often reveals that the organization never captured the identity, version, action history or decision context needed to support the assurance claim.

CISO operating questions / actions:

- Design control mechanisms to emit traceable records.
- Link change/configuration events to owners and approvals.
- Preserve provider notices and security-relevant service changes.
- Make incident and containment actions reconstructable.
- Store evidence with access, retention and integrity appropriate to its sensitivity.

Evidence-by-design does not mean logging everything. It means identifying which questions the organization should later be able to answer.

---

47. Separate Implementation Records from Defensible Evidence

Implementation teams produce policies, configuration records, architecture diagrams, test results, tickets, logs and supplier documents. These are candidate evidence sources, not automatic proof of effectiveness.

CISO operating questions / actions:

- Confirm the artifact relates to the actual scope and period.
- Confirm provenance and integrity are sufficiently understood.
- Identify whether it shows design, implementation, operation, monitoring or corrective action.
- Keep contrary evidence and failed tests rather than curating only favorable artifacts.

COM-015 owns the evidence-pack preparation layer. COM-018 may use organization-side information about what evidence exists, is stale, limited, contradictory or inaccessible for

CISO decision-making, but it does not create an evidence-sufficiency state, confidence score or assessment conclusion.

---

48. Preserve Scope, Period, Version and Provenance

Evidence loses assurance value when the organization cannot connect it to the system state it is supposed to support. The CISO should therefore expect the evidence system to preserve identity and time context.

CISO operating questions / actions:

- System/scope element.
- Owner/custodian.
- Period or event time.
- Model/service/configuration version where material.
- Source and location.
- Integrity/provenance information.
- Known limitations.

Later remediation should never be represented as if it operated throughout an earlier evidence or incident period.

---

49. Separate Provider Representations from Enterprise Observation

Provider evidence is often necessary, but it should retain its source identity. A provider statement, SOC report, security document or support response is not the same as the enterprise's own observation.

CISO operating questions / actions:

- Use the COM-015 evidence-source states consistently: **Customer-observed / Provider-provided / Shared-interface / Provider-only-inaccessible / Unknown**.
- Label provider-provided evidence explicitly and preserve its scope, period, criteria and limitations.
- Identify what remains provider-only or inaccessible; treat **Unknown** as an evidence gap until responsibility/evidence access is clarified.
- Corroborate material claims with enterprise evidence where feasible.
- Escalate evidence-access limitations when they affect high-consequence decisions.

These labels are evidence-collection aids, not control ratings. This distinction prevents outsourced operation from becoming outsourced truth, and a provider certificate or assurance report does not automatically transfer conformity to the customer's end-to-end AI system.

---

50. Preserve Unknowns, Contradictions and Negative Evidence

Credible assurance is not a collection of positive artifacts. The CISO should expect the record to include unresolved facts, failed tests, contradictory provider statements and evidence gaps where they are material.

CISO operating questions / actions:

- Confirmed / directly supported facts.
- Provider representations.
- Inferences or working hypotheses.
- Unknown or inaccessible elements.
- Specifically checked-but-not-observed conditions where the search basis matters.
- Contradictory evidence.
- Failed/negative tests and unresolved anomalies.

These are CISO evidence-management categories, not new UAIF enums or controlled assessment confidence states. Unknown does not mean failed; it means the supportable assurance boundary is narrower than it would be with stronger evidence.

---

51. Protect Sensitive Evidence

AI evidence can contain prompts, user content, source data, credentials, model details, employee information, privileged communications or security-sensitive architecture. Evidence collection therefore needs its own security and governance.

CISO operating questions / actions:

- Minimize sensitive content to what is necessary.
- Use controlled references rather than copying full data into general tickets.
- Apply access controls and retention appropriate to the record.
- Preserve legal privilege and investigation boundaries where applicable.
- Account for cross-border and provider contractual constraints.

An assurance programme that creates uncontrolled copies of sensitive evidence can increase the very risk it is meant to manage.

---

52. Keep Evidence Current Through Change

Evidence is only current relative to the system state it represents. A provider, model, permission, tool, data-source or architecture change can make previously strong evidence irrelevant to the new state.

CISO operating questions / actions:

- Link evidence to material change triggers.
- Identify which evidence classes should be refreshed or reconsidered after the change, according to materiality and the applicable evidence process.
- Preserve the historical record rather than overwriting it.

- Record what changed and which prior claims no longer apply.

Do not refresh evidence by changing a date while leaving an obsolete underlying artifact unchanged.

---

53. Handoff to COM-015 — Building an AI Assurance Evidence Pack

COM-015 owns the practitioner method for turning candidate records into a versioned evidence pack with indexing, provenance, source limitations and contradictions visible. COM-018 should consume that structure rather than duplicate it.

CISO operating questions / actions:

- Use COM-018 to determine what executive/security questions need evidence.
- Use COM-017 to generate implementation records during operation.
- Use COM-015 to prepare those records into an assurance evidence package.
- Use the resulting evidence position in CISO review and reporting.

Evidence-pack preparation does not itself create conformity, readiness or certification status.

---

54. Handoff Beyond Evidence Preparation

COM-015 governs public evidence preparation. The public practitioner journey ends at that evidence-preparation handoff. Any later assessment activity is separately governed and authorized. COM-018 does not expose controlled assessment methodology, participation procedures, scoring, sampling, finding logic, certification logic or assessor workflow.

CISO operating questions / actions:

- Ensure scope, owners and evidence custodians are available.
- Ensure material limitations, changes and provider evidence constraints are disclosed rather than hidden.
- Preserve the assessed/evidence period so later remediation is not rewritten as though it operated throughout an earlier period.
- Receive bounded results and route remediation/risk decisions through the correct organizational authority.
- Do not reverse-engineer assessor scoring, confidence, sampling, findings or certification logic from the public guide.

The controlled assessment-methodology family remains the authority for assessment work, and publication of practitioner guides does not by itself activate live assessment authority or certification availability.

---

55. CISO Claims Firewall

Security leaders are frequently asked for shorthand labels: secure, compliant, certified, approved, audit-ready. COM-018 should help the CISO resist claims that are broader than the supporting evidence and authority.

CISO operating questions / actions:

- Prefer scoped factual statements: what is implemented, observed, assessed or mapped.
- Identify the scope, date/period and material limitations.
- Do not use GAISSF certified/compliant/approved/endorsed language without the separately applicable authority.
- Do not convert security assurance into legal compliance or product/safety approval.
- Do not promise absence of future incidents or harm.

Truthful claims discipline is itself an assurance control because overstatement can create governance, legal and reputational risk.

Part VI — Design for AI Incident Readiness and Resilience

Use this Part when: the CISO needs to know whether material AI systems can be identified, reconstructed, contained, escalated and recovered before an incident occurs.

Working output: system-level readiness expectations and a clean handoff to UAIF™ / AI-IRF™ during active incident work.

Part checkpoint: responders know the system/version, evidence sources, containment surfaces, provider route and physical-safety interface where applicable.

Engineer investigability, containment and recovery into AI systems before failure, while preserving UAIF™, AI-IRF™ and physical-safety boundaries.

56. Design Incident Readiness Before the Incident

The CISO should treat incident readiness as a design property of material AI systems. If responders cannot identify the system, reconstruct consequential actions or constrain authority, the organization may discover the assurance gap only after harm.

CISO operating questions / actions:

- System/version identification.
- Identity and delegated-authority traceability.
- Evidence preservation capability.
- Containment/revocation surfaces.
- Provider escalation.
- Rollback/recovery.
- Physical-safety interface where applicable.

COM-018 describes readiness expectations; active response belongs to the incident-response guides and processes.

57. Detection Is Not Automatically an Incident

AI environments generate anomalies, model failures, safety events, policy violations and ordinary software/security alerts. Treating every observation as a formal incident can be as damaging as ignoring material events.

CISO operating questions / actions:

- Preserve the initial observation and facts.
- Begin proportionate containment when justified by risk even if classification is incomplete.
- Keep provider statements and hypotheses distinct from confirmed facts.
- Use UAIF™ where structured AI incident classification becomes appropriate.

Response may begin before classification is complete. A local signal, observation or event-under-triage may remain in the organization's existing SOC/case workflow while facts are developed. COM-018 does **not** create a UAIF `Observation` record type, authorize unsupported UAIF field values or create a new severity system.

---

58. Make Attribution and Reconstruction Possible

AI incidents can be difficult to reconstruct because behavior depends on model/provider version, context, retrieval, tools, delegated credentials and state. The CISO should require enough telemetry and records to reconstruct material decisions and actions.

CISO operating questions / actions:

- Which user/agent initiated the chain?
- Which model/service/version was used?
- What context/retrieval/tool inputs materially influenced the action?
- What permissions/credentials were active?
- What external systems changed state?
- What provider/configuration changes occurred around the event?

The goal is sufficient reconstruction for the incident question, not perfect replay of every internal model computation.

---

59. Preserve Evidence Proportionately

Incident evidence should preserve what is necessary to understand cause, scope, actions and consequence without indiscriminately copying sensitive AI data.

CISO operating questions / actions:

- Preserve logs/configuration/version identifiers before they rotate.
- Capture provider notices/support exchanges.
- Preserve tool/action history and credential/identity context.

- Record collection limitations and unavailable data.
- Use protected repositories or references for highly sensitive content.

Universal prompt/output retention is not required by this guide and may be inappropriate in many systems.

---

60. Pre-Design Containment Surfaces

Containment is faster and safer when teams already know where authority can be reduced. AI systems often have multiple containment surfaces; disconnecting the entire service is only one option.

CISO operating questions / actions:

- User/session.
- Identity/token/credential.
- Tool/API/connector.
- Agent or orchestration path.
- Model/provider route.
- Feature/workflow.
- Data/retrieval source.
- Physical command path.

Containment should target the authority or exposure that creates risk. For physical AI, cyber containment should be coordinated through the relevant safety/operational authority where the action can alter physical risk.

---

61. Prepare Provider Escalation and Dependency Containment

A provider-dependent system may be difficult to investigate or contain using enterprise controls alone. The CISO should ensure provider escalation is operational rather than only contractual.

CISO operating questions / actions:

- Named support/security escalation channels.
- Evidence request process.
- Service/version/change information where available.
- Credential or tenant containment options.
- Alternative provider/model/fallback arrangements where justified.
- Documented behavior when provider evidence cannot be obtained.

Provider non-response or evidence opacity should become a visible assurance limitation and may alter the operating decision for high-consequence systems.

---

62. Prepare for Agentic-AI Incidents

Agent incidents may involve delegated credentials, tool calls, state/memory, transactions and multi-step action chains. Traditional endpoint or application logs may not capture the full authority path.

CISO operating questions / actions:

- Preserve agent/service identity.
- Capture tool definitions and material schema/metadata changes where feasible.
- Record delegated credentials and permission state.
- Preserve consequential tool/action history.
- Include memory/state needed to understand persistent behavior.
- Design rapid revocation of agent authority.

The CISO should expect tool and metadata changes to be treated as potential security-relevant changes, not assume approval of the server identity permanently approves its future behavior.

---

63. Prepare for Physical-AI Incidents

Physical-AI incidents can make conventional cyber containment unsafe. The CISO should establish in advance who has authority when a security action could alter movement, process state, operator visibility or safe-state behavior.

CISO operating questions / actions:

- Identify the safety/operational decision authority.
- Understand local safe/degraded behavior.
- Preserve sensor, mode, command, override and relevant physical-event evidence.
- Coordinate cyber isolation with operational consequence.
- Separate security closure from physical return-to-service authorization.

COM-018 does not determine safe-state criteria or certify the system for renewed operation.

---

64. Handoff to UAIF™ and AI-IRF™

COM-018 should preserve a clean incident boundary. UAIF™ provides structured AI incident classification; AI-IRF™ provides the structured response framework. Neither should be recreated in the CISO guide.

CISO operating questions / actions:

- Use COM-018 to ensure readiness capabilities exist before an event.
- Use UAIF™ to structure incident classification and uncertainty.
- Use AI-IRF™ for containment, response, recovery and learning.
- Allow response to begin while classification continues when the risk warrants action.

Legal reporting, regulatory notification and sector-specific incident obligations remain separate determinations.

Where structured UAIF data is exchanged, use the current **TCR-STD-003 / published UAIF machine-readable schema package** and the applicable controlled profile. Machine-readable UAIF and AIIS artifacts can support structured case enrichment, controlled vocabularies, playbook integration, routing and other interoperability patterns, but automation does not remove the need for human/organizational decision authority where the context requires it. COM-018 does not invent additional UAIF record types or field values.

---

65. Recovery, Revalidation and Learning

Recovery should restore controlled operation, not simply restore availability. The CISO should ensure the organization can explain what changed, what was fixed and which assurance assumptions should be revisited.

CISO operating questions / actions:

- Restore/rebuild from trusted state where appropriate.
- Validate affected controls and dependencies.
- Revisit credentials, providers and data sources implicated in the event.
- Record corrective action and material change.
- Feed lessons into portfolio-level common dependencies and related systems.

For physical AI, security recovery remains distinct from safety/operational return-to-service authority.

---

66. Exercise the Assurance System

Exercises should test whether the management system works under uncertainty. The most valuable findings often concern ownership, access, evidence and decision authority rather than technical detection alone.

CISO operating questions / actions:

- Can responders identify the exact affected AI service/version?
- Can they revoke agent/tool authority?
- Can they obtain provider support/evidence?
- Can they preserve necessary evidence without violating handling constraints?
- Can security and safety/operations coordinate physical containment?
- Can executives receive a bounded statement of facts and unknowns?

COM-018 does not prescribe a universal exercise cadence. Frequency should reflect consequence, change, organizational requirements and the operating context.

Part VII — Operate, Challenge and Improve the Assurance Programme

Use this Part when: controls exist but the assurance position risks becoming stale because of change, exceptions, providers, incidents or unresolved remediation.

Working output: a repeatable management rhythm, material-change review, challenge model, exception governance and remediation portfolio.

Part checkpoint: material changes and unresolved gaps cannot disappear silently between formal reviews.

Create the management rhythm, challenge function, change governance and remediation discipline needed to keep the assurance position current.

67. Establish an AI Assurance Operating Rhythm

The assurance position should be reviewed often enough to remain meaningful, but a universal monthly/quarterly cadence would be artificial. The CISO should define a rhythm around material changes, incidents, provider shifts, exceptions and portfolio decisions.

CISO operating questions / actions:

- Portfolio review.
- Provider/dependency review.
- Material-change review.
- Exception/remediation review.
- Evidence staleness/limitation review.
- Incident-readiness review.
- Board/risk reporting where applicable.

The rhythm can combine scheduled reviews with event-driven triggers.

---

68. Review Material Change

Material-change review determines whether a previously valid security assumption remains valid. It should be proportionate to the change rather than automatically reopening every control.

CISO operating questions / actions:

- Which boundary or assumption changed?
- Which controls/mechanisms depend on it?
- Which evidence no longer represents the current state?
- Does incident readiness or containment need to change?
- Does the board/risk statement need updating?

A material change is not automatically an adverse finding. It is a trigger to reconsider the assurance basis.

---

69. Monitor Providers and Shared Dependencies

Portfolio-level assurance should show when many AI systems depend on the same provider, identity platform, tool server, data service or model route. These common dependencies can create concentration risk.

CISO operating questions / actions:

- Number and consequence of dependent systems.
- Provider evidence/access limitations.
- Change notification quality.
- Fallback/exit feasibility.
- Open incidents or material security changes.
- Shared credentials, connectors or control planes.

The objective is management visibility, not a provider risk score detached from context.

---

70. Run the Remediation Portfolio

CISO leadership should manage remediation as a portfolio of security consequence and dependency, not as a queue ordered only by age. AI gaps can be amplified by delegated authority, provider concentration or physical consequence.

CISO operating questions / actions:

- Affected systems and consequence.
- Exposure and exploitability/likelihood where supportable.
- Common dependencies.
- Compensating controls and detection.
- Evidence gap versus demonstrated control failure.
- Owner, blocker and implementation verification.

COM-018 does not create a universal finding-severity algorithm or remediation deadline.

---

71. Maintain Independent Challenge

Assurance weakens when the same team implements a control, selects the evidence and declares the conclusion without challenge. The CISO should maintain an appropriate review layer that is sufficiently separate from implementation ownership.

CISO operating questions / actions:

- Security architecture/design review.
- Peer or second-line challenge.
- Internal assurance/audit where appropriate.
- Independent testing for material controls.
- Provider claim challenge and corroboration.

Independence is contextual. COM-018 does not define assessor-independence rules or accreditation requirements.

---

72. Govern Exceptions and Residual Risk

Residual risk should reach the authority that can genuinely accept its consequence. The CISO should make sure the decision includes evidence, limitations and alternatives rather than only a risk label.

CISO operating questions / actions:

- What remains unresolved?
- What evidence supports the current view?
- What could happen if the assumption is wrong?
- What compensating controls exist?
- What event would require immediate reconsideration?
- Who has authority to accept or reject the condition?

Security may recommend against operation even when another authorized executive ultimately accepts the business risk.

---

73. Scale Across Business Units and Regions

A scalable programme should standardize the language of scope, evidence, incidents, exceptions and claims while allowing implementation mechanisms to reflect local architecture and obligations.

CISO operating questions / actions:

- Common portfolio fields and evidence expectations.
- Shared control patterns and reference architectures.
- Local regulatory/privacy/safety overlays.
- Regional escalation routes.
- Central visibility into material exceptions and high-consequence systems.

Consistency should mean comparable assurance reasoning, not identical tooling everywhere.

---

74. Handle Lifecycle Transitions

AI security assurance should follow the system through acquisition, redesign, transfer and retirement. Transitions often invalidate ownership, evidence and provider assumptions.

CISO operating questions / actions:

- Acquisition/onboarding of an existing AI system.
- Provider/model replacement.

- Major redesign or autonomy expansion.
- Business-unit or regional transfer.
- Merger/divestiture.
- Decommissioning and data/credential retirement.

Lifecycle closure should include evidence/record retention, credential revocation, provider offboarding and removal of obsolete monitoring/authority paths.

Part VIII — Measure, Report and Defend the Assurance Position

Use this Part when: the CISO needs a dashboard, board report or defensible enterprise-level statement of the AI security position.

Working output: a multi-dimensional dashboard and bounded executive narrative based on scope, evidence, incidents, providers, exceptions and unknowns.

Part checkpoint: the dashboard supports decisions without collapsing the enterprise position into a single score or certification-like status.

Design decision-useful CISO metrics and executive reporting without replacing evidence with a simplistic enterprise AI-security score.

75. Measure What Helps a CISO Decide

Metrics are useful when they change a decision. The CISO should resist dashboards that create confidence by aggregation while hiding the conditions that matter.

CISO operating questions / actions:

- Measure visibility, ownership and change exposure.
- Measure control/implementation state without calling it conformity.
- Measure evidence currency and limitations.
- Measure incident readiness and resilience capability.
- Measure provider/dependency concentration.
- Measure exception/remediation exposure.

A metric should have an owner, definition, source, limitation and decision use. If no decision changes when the metric changes, reconsider why it exists.

---

76. Design the CISO AI Assurance Dashboard

The dashboard should be multi-dimensional and drillable. Its purpose is to reveal where management confidence is strong, weak or uncertain—not to produce one enterprise AI-security number.

CISO operating questions / actions:

- Estate/ownership coverage.
- High-consequence and physical-AI exposure.

- Control/implementation state.
- Evidence state and staleness.
- Incident readiness/resilience.
- Provider/shared-dependency exposure.
- Exceptions/remediation.
- Material changes and unknowns.

Use colors or status indicators only if the underlying definition and evidence remain visible. Unknown should be allowed as a first-class management state. A dashboard color or locally defined status does not become an ODA3 assessment result, control rating, maturity level, readiness category or certification status.

76.1 Dashboard design check

Before adding a metric or status, ask:

- What management question does this answer?
- Which system or population does it cover?
- What source produces it?
- What does it omit?
- Who owns the underlying condition?
- What decision changes if the indicator changes?

If those questions cannot be answered, the indicator is likely decorative rather than assurance-relevant.

---

77. Measure Estate and Ownership Coverage

The first dashboard question is whether the CISO knows what needs attention. Coverage measures should expose blind spots rather than reward optimistic inventory counts.

CISO operating questions / actions:

- Known material AI systems/use cases.
- Systems with named business and security owners.
- Systems with identified model/provider/version paths.
- High-consequence systems with explicit escalation routes.
- Unresolved shadow/unknown AI areas.

COM-018 does not prescribe a universal 100% threshold or declare a coverage percentage to be an assurance result.

---

78. Measure Control and Implementation State

Portfolio reporting can show where security mechanisms are planned, in progress, implemented, blocked or retired, provided those labels remain project/operational states.

CISO operating questions / actions:

- Applicable-control work not yet owned.

- Material implementation dependencies.
- Blocked or exception-backed mechanisms.
- Controls affected by recent change.
- Shared/provider-controlled mechanisms with limited visibility.

An organization-defined `Implemented` label is only a management statement about implementation work. It does not establish operating effectiveness, evidence sufficiency, conformity, assessment readiness, certification readiness or an ODA3 control determination.

---

79. Measure Evidence Quality and Currency

The CISO dashboard should make weak evidence visible. A large evidence repository is not necessarily a strong assurance position.

CISO operating questions / actions:

- Evidence missing for material controls.
- Evidence materially older than the current system state.
- Provider-only/inaccessible evidence.
- Unresolved provenance/integrity concerns.
- Contradictory or negative evidence.
- Material claims dependent on assumptions.

The public guide can surface these conditions; formal evidence sufficiency remains governed by the applicable evidence/assessment methodology.

---

80. Measure Incident Readiness and Resilience

Incident-readiness measures should reflect whether material systems can be investigated and constrained—not only whether a playbook exists.

CISO operating questions / actions:

- Known responder/owner path.
- Version/identity reconstruction capability.
- Evidence preservation capability.
- Containment/revocation options.
- Provider escalation path.
- Fallback/rollback/recovery capability.
- Physical safety interface where relevant.

COM-018 does not create universal response-time thresholds or MTTR requirements.

---

81. Measure Provider and Concentration Exposure

A CISO should see where assurance relies heavily on external services or a small number of common platforms.

CISO operating questions / actions:

- Critical provider count and dependent systems.
- Single points of provider/model dependency.
- Evidence-access limitations.
- Change-notification limitations.
- Fallback/exit constraints.
- Shared tool/server/connector dependencies.

Concentration is not automatically unacceptable. It is a decision factor that should be visible alongside consequence and fallback capability.

---

82. Measure Remediation and Exception Exposure

Remediation reporting should help the CISO distinguish accepted temporary exposure from uncontrolled drift.

CISO operating questions / actions:

- Open material remediation items.
- Exceptions without clear owners.
- Repeatedly renewed exceptions.
- Dependencies blocking closure.
- Compensating-control state.
- High-consequence systems with unresolved gaps.

Age can be informative, but COM-018 does not define universal aging buckets or severity-based deadlines.

---

83. Report to the Board and Risk Committee

Board reporting should answer a small number of enterprise questions with evidence and boundaries rather than technical detail. The CISO should report what is known, what changed, what matters and what decision is required.

CISO operating questions / actions:

- Which material AI systems are in the assurance portfolio and which remain uncertain?
- Where are the most consequential control/evidence gaps?
- Which providers or common dependencies create material concentration?
- Are material systems incident-ready and recoverable?
- What significant incidents, near misses or changes altered the assurance position?
- Which residual risks or exceptions require executive attention?

A board report should distinguish management representation from independently challenged evidence and make material unknowns visible. This directly supports COM-008's board principle—**ask for the evidence, not the policy**—without turning the board guide into an assessment instrument. Avoid certifying/compliance language unless separately supported and authorized.

83.1 A concise board-report structure

A practical CISO board update can be organized around six questions:

1. **Scope:** Which material AI systems are in the security-assurance portfolio, and what remains unknown?
2. **Change:** What materially changed since the prior report?
3. **Control / evidence:** Which material control or evidence conditions require executive attention?
4. **Incidents / resilience:** What material events occurred, and can affected systems be investigated, contained and recovered?
5. **Dependencies:** Which provider or shared-platform dependencies materially affect the assurance position?
6. **Decision:** What risk, funding, constraint or escalation decision is required from the board/risk authority?

This is a reporting pattern, not a mandatory board template or ODA3 assessment output.

---

84. Answer the Questions Raised in COM-008

COM-008 gives boards a concise set of assurance questions. COM-018 should equip the CISO to answer them from the enterprise management system rather than from ad hoc slide preparation.

CISO operating questions / actions:

- Which AI systems are within structured security review? → portfolio/inventory.
- What evidence supports control operation? → evidence architecture and COM-015 handoff.
- Who is accountable? → ownership/decision rights.
- Can incidents be classified and responded to? → readiness plus UAIF™ / AI-IRF™ handoff.
- Do systems create physical effects? → D9/PAI-SF™ portfolio visibility.

The CISO answer should be bounded. 'We do not yet know' is preferable to an unsupported enterprise-wide assurance claim.

---

85. Report Uncertainty Explicitly

Executive reporting often suppresses uncertainty because dashboards favor clean statuses. AI assurance should do the opposite for material unknowns.

CISO operating questions / actions:

- Known: directly supported facts/observations.
- Represented: provider or owner statements not independently observed.
- Inferred: reasoned conclusion based on available evidence.
- Unknown/inaccessible: information not currently available.
- Contradicted: material evidence points in different directions.

These labels are executive-management aids only. They are not controlled assessment confidence categories, UAIF record types or replacements for the more detailed incident-investigation distinctions used in COM-013.

---

86. Avoid False Precision and Unsupported Claims

Composite scores can hide the difference between an unknown provider change, a missing owner, a failed control and a stale evidence artifact. COM-018 therefore rejects a universal enterprise AI-assurance score.

CISO operating questions / actions:

- Do not average incomparable risks into a single maturity number.
- Do not call implementation status a conformity result.
- Do not turn a dashboard color into certification/readiness.
- Do not imply regulator equivalence or legal compliance from framework use.
- Do not state zero risk or guaranteed effectiveness.

Use precise underlying facts and scoped judgements instead of numerical certainty the evidence cannot support.

Part IX — CISO Decision Playbooks

Use this Part when: a concrete decision is in front of the CISO now.

Working output: a bounded decision record that identifies scope, evidence, authority, unknowns, escalation and next action.

Part checkpoint: the decision is attributable, scoped, evidence-aware and does not claim authority the CISO or ODA3 does not hold.

Provide reusable decision structures for the situations most likely to require CISO judgement or escalation.

87. Decision Playbook — New AI System or Major Use Case

Before a material AI system enters production, the CISO should be able to decide whether the security boundary and assurance prerequisites are understood well enough to proceed, constrain or escalate.

CISO operating questions / actions:

- What is the business use and consequence of failure?

- Who owns the system and security controls?
- Which models/providers/data/tools/identities are in scope?
- What applicable GAISSF baseline governs the security work?
- Who implements and verifies the material controls?
- What evidence will be generated in operation?
- Can the system be investigated and contained?
- What changes require renewed review?

The playbook is a decision aid, not a deployment certification or mandatory universal gate.

---

88. Decision Playbook — Third-Party AI Provider

Provider onboarding should focus on the operating assurance relationship, not only the procurement questionnaire.

CISO operating questions / actions:

- What exact service/model/endpoint is approved?
- What enterprise data or credentials can it access?
- What provider changes could alter the security posture?
- What evidence can the enterprise obtain?
- How are incidents escalated and investigated?
- What provider/subprocessor dependencies matter?
- What fallback, migration or exit path exists?
- What material limitations should be reported if provider evidence is unavailable?

A strong contract can improve assurance but does not itself prove the provider's current operating state.

---

89. Decision Playbook — Agentic AI

Agentic AI raises the CISO decision from 'can the model generate this?' to 'what may the system do, with whose authority, and how can that authority be constrained?'

CISO operating questions / actions:

- Identify agent identity and sponsoring principal.
- Define allowed tools/actions and prohibited authority.
- Scope credentials, transactions, destinations and delegation.
- Place approval/policy gates where consequence requires them.
- Protect memory/state and tool definitions.
- Log consequential actions sufficiently for reconstruction.
- Provide rapid revocation and containment.

Do not impose a universal human-approval requirement. Approval placement should reflect consequence, reversibility and the operating context.

---

90. Decision Playbook — RAG / Enterprise Knowledge

The CISO should treat RAG as an authorization and integrity architecture, not merely a model-quality feature.

CISO operating questions / actions:

- Which sources are allowed?
- Who approves ingestion/indexing?
- Can users retrieve only what they are authorized to access?
- Can untrusted content inject instructions or alter behavior?
- How are connector scopes governed?
- Can the organization identify the source/version behind a consequential response?
- What happens when a source or connector changes?

Evidence of retrieval quality is different from evidence that the access path is secure and governed.

---

91. Decision Playbook — Material Model, Tool or Provider Change

A material change should trigger a focused assurance impact review rather than an automatic full restart or an automatic approval.

CISO operating questions / actions:

- What changed and why?
- Which system boundary/assumption is affected?
- Which controls or dependencies rely on the old state?
- Which evidence has become stale?
- Does incident reconstruction or containment change?
- Does provider responsibility or visibility change?
- Does executive reporting need updating?

Record the impact decision even when the conclusion is that no additional action is required.

---

92. Decision Playbook — Suspected AI Security Incident

The first CISO decision is to preserve options: protect people/systems, retain material evidence and avoid prematurely fixing the narrative around an unproven cause.

CISO operating questions / actions:

- Establish the affected system/owner and known facts.
- Preserve volatile evidence proportionately.
- Constrain credentials/tools/authority when justified.
- Keep provider representations and hypotheses separate.
- Escalate to UAIF™ classification and AI-IRF™ response as appropriate.
- Coordinate physical containment with safety/operations where relevant.

- Provide executives a factual statement of knowns and unknowns.

Legal/regulatory notification decisions remain with the applicable authorized process.

---

93. Decision Playbook — High-Consequence or Physical AI

For physical or other high-consequence AI, the CISO should ensure the security decision is connected to the full consequence path and the right non-security authorities.

CISO operating questions / actions:

- Identify perception/sensing assumptions.
- Define autonomy and command boundaries.
- Identify local safe/degraded behavior.
- Map independent override and containment interfaces.
- Use PAI-SF™ for detailed physical-AI security assurance.
- Bring safety/operational authority into incident and recovery decisions.
- Keep security closure separate from return-to-service authorization.

COM-018 is not a product-safety, functional-safety or deployment-approval standard.

---

94. Decision Playbook — Assurance Request from Board, Audit, Customer or Regulator

External or executive requests often arrive as broad questions: ‘Are we compliant?’, ‘Is the system secure?’, ‘Can you certify this?’ The CISO should translate the request into a bounded factual proposition before responding.

CISO operating questions / actions:

- What exact system/scope is being asked about?
- What period/version does the answer cover?
- What evidence supports the statement?
- What limitations, provider representations or unknowns are material?
- Is there a separately valid assessment or certification record that actually applies to this scope and request? If not, do not imply one.
- Does another function own the legal/regulatory conclusion?

Prefer a narrower defensible answer over a broader status label unsupported by evidence or authority.

---

95. Decision Playbook — Continue, Constrain, Degrade, Suspend or Retire

The CISO may need to recommend a change in operating state when confidence is materially impaired. The choice should reflect consequence, reversibility, detection, alternatives and the authorities involved.

CISO operating questions / actions:

- **Continue:** the available security basis supports continued operation for the current bounded management decision, subject to other applicable authorities.
- **Constrain:** reduce access, authority, population, tools or use case.
- **Degrade:** use a lower-authority or reduced-function mode whose use is governed by the applicable operational and risk authorities.
- **Suspend:** stop the affected capability while preserving evidence and recovery options where appropriate.
- **Retire:** recommend removal of the capability when the security basis no longer supports continued operation, with the final decision made by the applicable organizational authorities.

These are illustrative management options, not an ODA3 operating-state taxonomy, validation status, assessment result, certification status or unilateral CISO authority over business, legal, safety or sector decisions.

Part X — Practitioner Tools, Boundaries and Publication Controls

Use this Part when: the organization needs reusable management artifacts or a final claims/authority check.

Working output: a charter, portfolio register and decision-rights/claims checklist that can be adapted to existing enterprise tooling.

Part checkpoint: the tools support management without becoming mandatory GAISSF artifacts, assessor workpapers, readiness instruments or certification gates.

Provide public-safe management tools, source boundaries, GEL attribution and working-publication controls.

96. CISO AI Security Assurance Charter — Practitioner Tool

The charter can be implemented as a short internal document. The following fields are public-safe and deliberately avoid assessment or certification language.

Best used when: AI security responsibilities are fragmented, the CISO is being asked to answer for systems outside the security function’s control, or incident/risk decisions reveal unclear authority.

Minimum usable version: one page is enough if it names the mandate, scope, decision rights, interfaces, exception/risk route and executive reporting path. Add detail only where it changes a real decision.

CISO operating questions / actions:

- Mandate/purpose.
- In-scope portfolio criteria.
- Executive sponsor and CISO authority.
- System/control ownership expectations.
- Interfaces with AI governance, engineering, privacy/legal, supplier risk, SOC/CSIRT, internal assurance and safety/OT.
- Exception/risk-acceptance route.
- Material-change triggers.
- Board/risk reporting route.
- Claims and evidence boundary.

Use the charter to align organizational authority. Completion of the template is not an ODA3 readiness designation.

---

97. AI Assurance Portfolio Register — Practitioner Tool

The portfolio register should help the CISO identify where attention is needed without converting the estate into a single risk score. Suggested fields can be adapted to organizational tooling.

Best used when: the CISO has more than a small number of material AI systems, multiple providers, shared platforms or recurring executive questions.

Minimum usable version: start with system/use case, business owner, security owner, model/provider, major authority/consequence characteristic, current material gap/unknown, last material change and next decision. Expand into the full field set as the portfolio matures.

Do not duplicate inventories unnecessarily: where CMDB, cloud, model registry, GRC or supplier systems already hold authoritative fields, link or federate them instead of creating a second uncontrolled source of truth.

CISO operating questions / actions:

- System/use case and business owner.
- Security/control owner.
- Model/provider/version path.
- Data/retrieval/tool/agent characteristics.
- Physical/high-consequence flag.
- Applicable control/profile reference.
- Implementation state and major dependencies.
- Evidence state/last material evidence period.
- Provider visibility limitations.
- Incident/change history.
- Exceptions/remediation.
- Material unknowns and next decision.

The register is a management inventory, not an assessment record or certification register.

---

98. CISO Decision Rights and Claims Checklist — Practitioner Tool

Before a high-consequence security decision or external assurance statement, the CISO can use a short combined checklist.

Best used when: the statement may reach the board, audit, a customer, a regulator, a safety/operational authority or another party that may rely on it.

Fast use: answer the checklist in writing when the decision is material. A short record of scope, evidence, unknowns and authority is more useful than a polished but untraceable conclusion.

CISO operating questions / actions:

- Do I have authority to make this decision, or only to recommend/escalate?
- Which system/version/scope does the decision cover?
- What evidence supports it and what remains represented/inferred/unknown?
- Which business, legal, privacy, safety or sector authorities also apply?
- Does a provider limitation materially narrow the conclusion?
- Am I using any certified/compliant/approved/endorsed wording without separately verified authority?
- Would material new evidence or change require the decision to be revisited?

The checklist is designed to slow down overclaiming, not to create a formal approval workflow.

---

99. Notably Absent, Source Traceability and Next Reading

COM-018 is intentionally bounded. Its value depends on making clear what a CISO guide cannot establish.

Notably Absent:

- no new GAISSF control requirement or CISO-specific conformity profile;
- no universal AI-security maturity/readiness score;
- no universal risk/severity formula or fixed assurance SLA;
- no controlled assessor scoring, sampling, finding or certification logic;
- no certification availability or assessor authorization created by this guide;
- no partner/empanelment status created by this guide;
- no regulator recognition or legal-compliance conclusion;
- no functional-safety, product-safety or return-to-service approval;
- no guarantee that incidents or harm will not occur.

99.1 Source hierarchy

For exact requirements, use the current authoritative source rather than this practitioner guide.

Source / family	Role in COM-018
GAISSF-NOR-001 — Framework Standard	Framework-level interpretation and authoritative GAISSF baseline.

Source / family	Role in COM-018
GAISSF-NOR-004 — Control Catalogue	Canonical control library and domain/control architecture.
Current GAISSF release / reconciliation records	Supersession, release-state and profile-reconciliation provenance.
GEL v1.0	Governing licence/use boundary for the GAISSF Ecosystem.
UAIF™ controlled specification / TCR-STD-003	Incident-classification architecture and published machine-readable schema companion.
AI-IRF™ controlled technical sources / AIIS v1.0.0	Incident-response architecture and published machine-readable control/playbook encoding.
PAI-SF™ canonical sources	Detailed physical-AI security-assurance requirements and physical/safety interface.
WHP practitioner series	Informative implementation, evidence, assessment-participation and incident/physical-AI guidance.

Historical source titles remain identifiable as historical source titles and are not silently rewritten into current market-facing terminology.

99.2 Current practitioner references

1. **The AI Assurance Imperative: A Board Member's Guide to GAISSF™** — ODA3-2026-08-WHP-COM-008 — board-level oversight questions that COM-018 helps management answer.
2. **When AI Breaks: The ODA3 Guide to AI Incident Response** — ODA3-2026-08-WHP-COM-011 — short incident-response orientation.
3. **Classifying AI Incidents: A Practical Guide to UAIF™** — ODA3-2026-08-WHP-COM-012 — UAIF™ classification and uncertainty discipline.
4. **AI Incident Response Field Guide: Using UAIF™ and AI-IRF™ for Real-World AI Incidents** — ODA3-2026-08-WHP-COM-013 — combined classification/response practitioner treatment.
5. **Securing Physical AI: A Practitioner's Guide to PAI-SF™** — ODA3-2026-08-WHP-COM-014 — detailed physical-AI security-assurance and safety interface.
6. **Building an AI Assurance Evidence Pack** — ODA3-2026-08-WHP-COM-015 — evidence indexing, provenance, limitations, contradictions and handoff preparation.
7. **Implementing GAISSF™** — ODA3-2026-08-WHP-COM-017 — detailed implementation and implementation-side verification.

For implementation-to-assessment continuity, the practitioner handoff remains:

scope / applicability → implementation → evidence generation → COM-015 evidence preparation → separately governed assessment activity, where independently authorized

This is a functional relationship, not a mandatory chronological workflow.

99.3 Machine-readable companions and interoperability wording

The principal machine-readable companions relevant to COM-018 include:

- **TCR-STD-003 / the published UAIF machine-readable schema package**; and

- **AIIS v1.0.0**, the published AI-IRF machine-readable control/playbook encoding.

Future interoperability work may include product integrations, field mappings, connectors, adapters, SDKs, additional profiles, reference integrations or validation/conformance tooling. Such work should not be described in a way that implies the core machine-readable schemas are still awaiting publication.

---

100. Publication Metadata, GEL Attribution and Final Publication Statement

100.1 Publication metadata

- **Document ID:** ODA3-2026-08-WHP-COM-018
- **Title:** The CISO's Guide to AI Security Assurance
- **Subtitle:** A Decision Guide to Controls, Evidence, Incident Readiness and Executive Accountability
- **Document Type:** WHP — Executive / Practitioner Decision Guide
- **Framework Context:** GAISSF Ecosystem / Operational AI Assurance
- **Classification:** Public — GEL v1.0
- **Publisher:** ODA3 Institute
- **Publication date:** 19 August 2026 (supersedes R9, 13 August 2026)
- **Publication state:** CORRECTED FINAL (R10)
- **Copyright:** © 2026 ODA3 Pvt Ltd

Suggested citation: ODA3 Institute (2026), *The CISO's Guide to AI Security Assurance: A Decision Guide to Controls, Evidence, Incident Readiness and Executive Accountability*, ODA3-2026-08-WHP-COM-018 (R10).

No DOI, ISBN, accreditation number, certification number or regulator-recognition identifier is assigned by this guide.

100.2 Naming and marks

Current public practitioner usage in COM-018 is:

- **GAISSF™ — Global AI Safety and Security Framework**
- **UAIF™ — Unified AI Incident Framework**
- **AI-IRF™ — AI Incident Response Framework**
- **PAI-SF™ — Physical AI Security Framework**
- **GAISSF Ecosystem** — umbrella term for the integrated ODA3 framework ecosystem

Historical source titles may preserve older naming or mark treatment for source identification only.

100.3 GEL attribution and use boundary

GAISSF™, UAIF™, AI-IRF™, and PAI-SF™ are frameworks of ODA3 Institute and are referenced under the GAISSF Ecosystem License (GEL) v1.0.

GEL v1.0 is the governing licence for use of the GAISSF Ecosystem. Public availability and Internal Use do not create unrestricted third-party Commercial Use rights.

External professional delivery—including consulting, implementation, integration, assessment, audit, paid training, managed services, software/SaaS embedding, certification or other client delivery—remains subject to the applicable GEL commercial rights and any separately applicable ODA3 programme authorization.

Factual alignment or interoperability statements, where accurate and otherwise permitted, do not themselves imply ODA3 certification, endorsement, audit, approval, partner status or official authorization.

100.4 Final Publication Statement

This document is the released **CORRECTED FINAL (R10)** edition of **ODA3-2026-08-WHP-COM-018 — The CISO's Guide to AI Security Assurance**, superseding the R9 FINAL edition of 13 August 2026.

R10 is a controlled post-release reconciliation: it removes public references to ODA3-2026-08-WHP-COM-016 (withdrawn from public publication effective 18 August 2026) and re-pins UAIF™ and AI-IRF™ framework maturity language to their current Final Publication v1.0 status. No change was made to the 10-Part / 100-section architecture, decision playbooks, practitioner tools, evidence architecture, dashboard logic or claims firewall.

CORRECTED FINAL status applies to this executive/practitioner guide only. It does not create or alter the maturity, authority, availability, recognition or approval state of any underlying framework, assessment methodology, certification programme, partner programme or external authority.

The publication state of this guide does not create or alter:

- framework maturity;
- live assessment authority;
- certification availability;
- assessor authorization;
- partner or empanelment status;
- regulator recognition;
- legal-compliance status;
- product or sector approval;
- functional-safety approval;
- physical return-to-service authority; or
- commercial-delivery rights.