FIOR

The Enforcement Layer for AI

Automated Compliance with MAS Guidelines on AI Risk Management for Financial Institutions

Author: FIOR Group

Date published: 17th November 2026

Version: 1.7

Classification: Commercial in Confidence



Executive Summary

FIOR provides a deterministic, self sovereign identity, governance and enforcement layer for AI agents. FIOR controls agents before they can act.

On 7 October 2026 the Monetary Authority of Singapore (MAS) issued its final Guidelines on Artificial Intelligence Risk Management (AIRG), following Consultation Paper P017-2025 of November 2025. The Guidelines apply to all financial institutions (FIs) in a proportionate manner and set out MAS' expectations on AI oversight, key AI risk management systems, policies and procedures, AI life cycle controls, and the capabilities and capacity needed for the use of AI. They take effect on 7 October 2027: FIs may meet the expectations in Sections 3 and 4 from that date, and those in Sections 5 and 6 by 7 October 2028.

FIOR AI Gateway gives FIs the technical controls the Guidelines call for — AI identification (including shadow AI), guardrails, human oversight, cybersecurity, monitoring and change control — together with the evidence their reviewers will ask for. It runs within the FI's own environment, on Kubernetes (including air-gapped installations), as a dedicated appliance or on AWS, and enforces policies the FI sets to its own risk appetite, well ahead of the Guidelines' effective date.

The FIOR product is self sovereign, meaning no data is required to leave the enteprise boundary in return for identity and nothing is stored on external databases. The method of identity uses NIST approved post quantum cryptograhy making it highly secure and quantum resistant.

> You can control your agents before the act with full confidence that the authority you give them is secure and cannot be spoofed or tampered with.
> *David Williams, Founder, FIOR*

This document maps FIOR's capabilities against the final Guidelines, paragraph by paragraph. For each expectation it states what FIOR enforces, what FIOR evidences, and what remains the FI's own governance decision.

FIOR has patent-pending technology for the identity and enforcement layer for AI agents: X.509 agent identity, deterministic policy enforcement, tamper-evident audit and rapid, fleet-wide revocation, designed for regulated and sovereign environments. FIOR's security and differentiation have been recognised in several ways since it launched its first software in May 2026.

© 2026 FIOR Group Limited
Visit our website: fior.group

FIOR AI Gateway achieved the status “NCSC CRTF Cyber Resilience Assessed” under the UK National Cyber Security Centre’s Principle Based Assurance (PBA) scheme (evaluation PBA-004, report issued 9 November 2026). All 72 of the security claims FIOR makes for the product were assessed as met (https://fior.group/resources/) .

IDC published a Vendor Profile on FIOR in May 2026. It said:

> *“FIOR is addressing emerging changes in AI agent security and governance through its patented identity infrastructure and enforcement-centric approach. Unlike traditional identity and access management (IAM) or AI security platforms that focus on detection or compliance reporting, FIOR enforces cryptographic identity and policy at the network boundary and agent level. Its Gateway and Auth products deliver cryptographically anchored authentication and real-time policy enforcement for AI agents, machines, and IoT, addressing the emerging risks of agentic automation in critical sectors. FIOR's focus on trust enforcement, rapid revocation, and audit-grade governance positions it as a differentiated player in the rapidly evolving AI security landscape. This reflects a broader shift in AI security toward deterministic control of autonomous systems, rather than retrospective visibility into their behaviour.”*

© 2026 FIOR Group Limited
Visit our website: fior.group

Table of Contents

Executive Summary *2*

1. How to read this document *5*

Responsibility legend 5

2. MAS AIRG at a glance *6*

3. FIOR platform architecture *7*

Identity and integrity 7

4. Responsibility summary *7*

5. Detailed mapping *10*

1 Introduction 10

2 Application of the Guidelines 11

3 AI Oversight 12

4 Key AI Risk Management Systems, Policies & Procedures 13

5 AI Life Cycle Controls 16

6 AI Capability & Capacity 24

6. Evidence artefacts *25*

7. Deployment options and rollout *25*

Typical 12-week rollout 25

8. Appendix A — Glossary *27*

9. Appendix B — Gateway policy controls reference *28*

© 2026 FIOR Group Limited
Visit our website: fior.group

1. How to read this document

This document looks at the MAS Guidelines on Artificial Intelligence Risk Management (7 October 2026) paragraph by paragraph and maps each expectation to the FIOR capability that supports it. It is written for two audiences:

- **Compliance, risk and audit teams —** use the mappings in §4–§5 as a working control matrix, and §6 for the evidence each control produces.
- **Architects and engineers —** use §3 and §7 to explain the architecture and deployment.

Scope. The mapping describes the FIOR AI Gateway and its customer console as released in November 2026, against the final Guidelines issued by MAS on 7 October 2026. Paragraph numbers follow the final Guidelines.

Responsibility legend

Responsibility	Meaning
FIOR control	FIOR operates the control. The FI enables and configures it; FIOR enforces it and records the evidence.
Shared	FIOR provides the enforcing control or system of record for part of the expectation; the FI completes it with its own configuration, decisions or documentation.
Customer governance	A governance responsibility that sits with the FI's Board, senior management or control functions. FIOR supplies operational records that support it.

© 2026 FIOR Group Limited
Visit our website: fior.group

2. MAS AIRG at a glance

The Guidelines are structured in six sections and are applied proportionately to the FI's size and risk profile, the nature of its activities, its use of AI and the risk materiality of each AI use case. An FI whose AI use is unlikely to have a material adverse impact may instead apply basic AI governance policies (paragraph 2.5).

Section	Paragraphs	Focus
§1 Introduction	1.1 – 1.12	Scope, definition of AI, related frameworks, effective dates, and risks from AI, Generative AI and AI agents (including shadow AI).
§2 Application of the Guidelines	2.1 – 2.5	Proportionate application; criteria for, and minimum content of, basic AI governance policies.
§3 AI Oversight	3.1 – 3.6	Board and senior management responsibilities; coordination across functions; local senior management accountability under group frameworks.
§4 Key AI Risk Management Systems, Policies & Procedures	4.1 – 4.13	AI identification (including third-party and shadow AI), AI inventory, risk materiality assessment (impact, complexity, reliance).
§5 AI Life Cycle Controls	5.1 – 5.26	Data, transparency, fairness, human oversight, third-party AI, selection, evaluation, cyber, reproducibility, pre-deployment review, monitoring, change management, decommissioning.
§6 AI Capability & Capacity	6.1 – 6.3	Personnel competence, capability review, technology infrastructure.

The Guidelines take effect on 7 October 2027; FIs may meet the expectations in Sections 3 and 4 from that date and in Sections 5 and 6 by 7 October 2028. They complement existing MAS legislation and guidelines, including the FEAT principles, the Technology Risk Management Guidelines and MAS' expectations on outsourcing and third-party services.

© 2026 FIOR Group Limited
Visit our website: fior.group

3. FIOR platform architecture

FIOR runs as an inline control plane between AI agents and the AI services, models and tools they use. Every request is authorised individually against the agent's identity and the FI's policies before it leaves the FI's environment, and every decision is recorded. Identity is provided within a Self-Sovereign and quantum resistant platfrom inside the enterprise boundary. No date on the existence or activity of an agents needs to be sent to any external identity provider.

Component	Role	AIRG paragraphs primarily served
FIOR Gateway (policy engine)	Authenticates each agent identity, evaluates the FI's policies on every request with default deny, applies revocation and quarantine, and records each decision with its reason and the deciding policy.	1.10, 1.11, 4.2, 5.3, 5.15, 5.16, 5.23
FIOR Proxy	Agent-side enforcement point: mutual-TLS listener, inline content controls (credential detection, PII masking, threat inspection), tool-call governance and human-approval holds, rate and cost limits.	1.10, 1.11, 5.4, 5.9, 5.15, 5.22(b)
FIOR Sensor	Passive network sensor that discovers AI applications in use and flags any not matched to an enrolled FIOR identity.	1.10(f), 4.2, 4.4
Customer console	Observe · Control · Report: event stream and drill-down, agent and certificate inventory, policies with version history, approvals, revocation, audit trail, cost reporting and responsible-officer review.	3.x, 4.x, 5.9, 5.23 – 5.26

Identity and integrity

- **Self-sovereign agent identity.** Each agent holds an X.509 identity issued under the FI's own certificate authority chain. There is no third-party identity provider, and no external service sits in the authentication path. In production, FIOR generates its signing CA key inside the FI's deployment and the FI's own root CA — which can be held in the FI's HSM or a managed CA service — certifies it through a standard certificate signing request, so no CA private key is ever exported or pasted. In new deployments, agent identity is accepted only by mutual TLS: on every request the client certificate must chain to the FI's root and match the agent's record (status, expiry, fingerprint and permitted models), with replay protection. The FI can additionally require per-request signatures and post-quantum signatures by policy.
- **Post-quantum readiness.** Agent credentials support hybrid Ed25519 + ML-DSA-65 (FIPS 204) signatures; the proxy offers hybrid X25519 + ML-KEM-768 key exchange on its outbound TLS connections.

© 2026 FIOR Group Limited
Visit our website: fior.group

- **Tamper-evident audit.** The administrative audit trail is a server-side SHA-256 hash chain with an immutability control and an in-console verification. Gateway decision evidence is protected by a database-enforced minimum retention period.
- **Independently verifiable provenance.** An optional provenance mode records gateway events in a Merkle-tree log with checkpoints signed using Ed25519 and ML-DSA-65 and RFC 3161 timestamps, exportable and verifiable offline with the fior-ctl tool.

© 2026 FIOR Group Limited
Visit our website: fior.group

4. Responsibility summary

This mapping covers 74 paragraphs and sub-paragraphs of the final Guidelines. The table below shows, for each section, how responsibility divides between FIOR and the FI.

Section	Paragraphs	FIOR control	Shared	Customer governance
1 Introduction	4	2	2	0
2 Application of the Guidelines	3	0	2	1
3 AI Oversight	6	0	3	3
4 Key AI Risk Management Systems, Policies & Procedures	13	0	9	4
5 AI Life Cycle Controls	45	12	15	18
6 AI Capability & Capacity	3	0	1	2
Total	**74**	**14**	**32**	**28**

FIOR's enforcing controls are concentrated where MAS expects technical controls — Generative AI and agent risks, shadow AI, identification, guardrails, cybersecurity, human-oversight design, monitoring and change control. Board, senior-management and model-development responsibilities remain with the FI, supported by FIOR's records.

© 2026 FIOR Group Limited
Visit our website: fior.group

5. Detailed mapping

Each paragraph of the final Guidelines in scope is listed with the FIOR capability, the responsibility allocation and the evidence an FI or reviewer would inspect. MAS text is summarised; refer to the Guidelines for the verbatim wording.

1 Introduction

1.1 – 1.2 — Scope, related frameworks and FEAT [Shared]	
MAS expectation	The Guidelines set supervisory expectations for AI risk management and complement existing MAS legislation and guidelines; FEAT principles continue to apply; applicable industry standards should be taken into account.
FIOR support	FIOR provides an enforcement and evidence layer for AI agents and the AI services they call. The FI's AI risk framework, FEAT practices and governance artefacts remain under its ownership; FIOR supplies the operational controls and records that evidence them.
Evidence	Cross-cutting

1.3 – 1.4 — Definition of AI [Shared]	
MAS expectation	AI covers machine-based systems or models that derive outputs through learned premises, including Generative AI and AI agents (including multi-agent systems); rule-based tools are excluded.
FIOR support	Every AI agent operating through FIOR holds a distinct cryptographic identity recording its provider, operator, permitted models, agent harness and host operating system — a system-level record of each AI system in use. Classification of use cases remains with the FI.
Evidence	Control → Agents · certificate records

1.10 — Generative AI risks [FIOR control]	
MAS expectation	Generative AI amplifies security, privacy, legal/IP, third-party, operational and human-factor risks — the last now expressly including shadow AI and automation bias.
FIOR support	Security: inline threat inspection of prompts with MITRE ATT&CK-mapped findings, configured to fail closed in the FIOR appliance. Privacy: credential detection, PII masking (fail closed) and log redaction. Third-party: allow-listed upstream hosts and regions with certificate pinning, and centrally managed model routes. Operational: rate, token, concurrency and cost-anomaly limits. Human-factor: FIOR Sensor detects shadow AI, and designated agent actions are held for human approval. Every decision is recorded with its reason.
Evidence	Observe → Events · Observe → Observations

© 2026 FIOR Group Limited
Visit our website: fior.group

1.11 — AI agent risks [FIOR control]	
MAS expectation	Agents with autonomy and tool access may execute unauthorised or erroneous actions, or be used to exfiltrate data or execute malicious commands at scale.
FIOR support	FIOR is built for autonomous agents: a cryptographic identity per agent; tool-call governance (allow and deny lists, argument and path restrictions, call quotas, loop detection, schema and pinned-contract checks, and human approval for designated tool calls), applied before the agent receives the call; fleet-wide quarantine of an agent on threat detection; and revocation.
Evidence	Control → Agents · Control → Approvals · Observe → Events

2 Application of the Guidelines

2.2 — Full AI risk management framework [Shared]	
MAS expectation	Unless basic policies suffice, establish oversight frameworks with AI identification, risk materiality assessment and an AI inventory (Sections 3–4), and apply relevant life cycle controls, capabilities and infrastructure (Sections 5–6).
FIOR support	FIOR supplies the system-backed elements of the framework: AI identification through FIOR identities and network discovery, a live agent inventory, and enforced life cycle controls, as mapped in Sections 3–6 below. The framework itself is the FI's.
Evidence	Sections 3 – 6 of this mapping

2.3 – 2.4 — Eligibility for basic AI governance policies [Customer governance]	
MAS expectation	Basic policies may be applied where poor performance or unavailability of the AI used is unlikely to have a material adverse impact on the FI, its customers or other stakeholders.
FIOR support	The eligibility assessment is the FI's. FIOR's records of which AI services are used, by which agents, for what volume of activity and with what tool access give the FI an evidence base for it, and for showing that use remains assistive.
Evidence	Observe → Events · Report → Cost

2.5 — Basic AI governance policies [Shared]	
MAS expectation	Clear accountability; permitted and prohibited uses with human review where required; an approved list of AI tools and an approval process for new ones; staff education; regular compliance checks; periodic and trigger-based review.
FIOR support	(a) Responsible officers can be designated and assigned in the console. (b) Permitted and prohibited uses become enforceable gateway policy: approved providers, models, domains and tools are permitted and everything else is refused by default; prompts containing credentials can be refused and personal data masked before it reaches an external AI service; designated agent actions are held for human approval. (c) The approved list of AI tools is enforced, not just published; adding one is a policy change that can require dual-control approval. (e) Decision records and shadow-AI discovery support compliance checks. (d) and (f) are FI processes.
Evidence	Control → Policies · Control → Approvals · Observe → Observations

© 2026 FIOR Group Limited
Visit our website: fior.group

3 AI Oversight

3.1 — Board and senior management role [Customer governance]	
MAS expectation	The board and senior management establish and oversee robust AI risk management frameworks, policies and procedures.
FIOR support	FIOR supplies the operating evidence the board and senior management rely on: enforced policies with version history, the tamper-evident audit trail, and responsible-officer review records with CSV/PDF export.
Evidence	Report → Audit · Corporate → Officer review export

3.2 — Existing risk frameworks and oversight [Shared]	
MAS expectation	Ensure existing risk frameworks identify, assess and address AI risks; maintain effective oversight; ensure AI use does not conflict with regulatory requirements.
FIOR support	The gateway enforces quantitative limits — request rates, token budgets, cost anomalies, concurrency and error thresholds — and records every decision, giving measurable indicators and evidence of adherence. Refusals, quarantines and unregistered-AI alerts are recorded for escalation. Risk identification and appetite are FI responsibilities.
Evidence	Observe → Events · Report → Cost · Observe → Observations

3.3 — Consistent standards and coordination [Customer governance]	
MAS expectation	Consistent standards, clear accountability and coordinated management of AI risk across functions, for example through cross-functional committees.
FIOR support	Coordination structures are the FI's. FIOR gives every function one shared, consistent record — the same policies, decisions and audit trail — as the evidence base for cross-functional oversight.
Evidence	Report → Audit · Officer review export

3.4 — Board responsibilities [Customer governance]	
MAS expectation	Approve and review the governance approach and strategic direction for AI; address AI risks in the risk appetite framework with qualitative statements and quantitative measures; set roles; ensure AI understanding; review regularly.
FIOR support	Quantitative risk appetite measures approved by the board can be implemented directly as enforced gateway limits, and prohibited use categories as refusals, so approved appetite and operational enforcement stay aligned. Gateway decision records support measures such as the number of AI-related refusals and interventions.
Evidence	Control → Policies · Observe → Events

© 2026 FIOR Group Limited
Visit our website: fior.group

3.5 — Senior management responsibilities [Shared]	
MAS expectation	Implement and review AI risk frameworks; institute controls across the AI life cycle; set roles, including control functions; escalate material risks and AI incidents; update the board; resource competence.
FIOR support	(a) Policies are enforced at the gateway on every request, with version history and change notes. (c) FIOR's controls operate across each agent's operational life cycle. (d) Role-based access is enforced in the database. (e) Approval requests, responsible-officer assignment and delegation, revocation and quarantine provide an escalation and action path; dual control requires a second authorised person to approve policy changes. (f) Audit and officer-review exports support board reporting. (b) and (g) are FI responsibilities.
Evidence	Control → Policies · Control → Approvals · Corporate → Officer assignments

3.6 — Local senior management under group frameworks [Shared]	
MAS expectation	Local senior management remains accountable where group frameworks are used, with visibility, timely reports and data, clear escalation paths, and the ability to demonstrate oversight to MAS.
FIOR support	FIOR runs within the FI's own environment, so local management has direct access to Singapore decision records, approvals and the audit trail, and can export them to demonstrate oversight. The escalation design is the FI's.
Evidence	Observe → Events · Report → Audit · Control → Approvals

4 Key AI Risk Management Systems, Policies & Procedures

4.1 — Identification, inventory and materiality [Shared]	
MAS expectation	Consistent, robust approaches to identify, inventorise and assess AI, with proportionate controls.
FIOR support	Identification and inventory of AI agents are system-backed by FIOR identities and network discovery; controls are applied per agent in proportion to the FI's assessment. The materiality methodology is the FI's.
Evidence	Control → Agents · Control → Sensors

4.2 — Consistent AI identification [Shared]	
MAS expectation	Systems, policies and procedures for consistent identification of AI use, supported by tools and systems, including AI in services from material third-party providers.
FIOR support	With default deny, AI traffic through the gateway must present a FIOR identity to be served. FIOR Sensor passively observes network traffic to AI services and raises an alert for any AI application not matched to an enrolled identity, which a reviewer then confirms or dismisses. Identification of AI embedded within third-party services is completed through the FI's third-party processes.
Evidence	Observe → Observations · Control → Sensors

© 2026 FIOR Group Limited
Visit our website: fior.group

4.3 — Roles for identification [Shared]	
MAS expectation	A designated control function with independent oversight, acting as final arbiter; documentation and regular review.
FIOR support	Role-based access (provider administrator, corporate security, operator, responsible officer, super administrator) is enforced in the database. Review decisions on discovered AI are recorded. Designation of the control function is an FI decision.
Evidence	Admin → Role matrix · Report → Audit

4.4 — Unidentified AI and shadow AI [Shared]	
MAS expectation	Keep residual risk from unidentified AI, including shadow AI, within risk appetite, using measures such as network monitoring and data loss prevention to detect unauthorised use.
FIOR support	FIOR provides the technical measures MAS names. Network monitoring: FIOR Sensor detects AI applications in use that are not enrolled. Data loss prevention: credential detection and PII masking on AI traffic through the gateway. Default deny refuses any AI request no policy permits. Setting the risk appetite is the FI's.
Evidence	Observe → Observations · Control → Policies

4.5 — Accurate AI inventory [Shared]	
MAS expectation	Maintain an accurate inventory, with an appropriate update frequency, reflecting new, updated and decommissioned AI.
FIOR support	The identity inventory updates continuously: issuance, expiry and revocation change an agent's status automatically, so a decommissioned agent is reflected as soon as it is revoked.
Evidence	Control → Agents · Provider → Revocation

4.6 — Linkage to other inventories [Shared]	
MAS expectation	Clear linkages between the AI inventory and other inventories, e.g. through consistent identifiers.
FIOR support	Each agent carries a stable, unique identifier that the FI can use as the consistent identifier linking its AI inventory, data inventories and third-party registers.
Evidence	Control → Agents

4.7 — Inventory attributes [Shared]	
MAS expectation	Purpose, scope of use, model type, data, dependencies, life cycle status, materiality rating, review status, owners, documentation; for AI agents, agent-specific identifiers, the tools and systems the agent can access, and its guardrails.
FIOR support	FIOR records the agent-specific attributes MAS highlights. Each identity holds agent name, provider, operator, permitted models, permitted domains, agent harness, operating system, issue and expiry dates and status. The tools and systems an agent may reach and the guardrails applied to it are defined in its enforced policies. Purpose, materiality rating and review status are kept in the FI's AI register against the FIOR identifier.
Evidence	Control → Agents · Observe → Identity · Control → Policies

© 2026 FIOR Group Limited
Visit our website: fior.group

4.8 — Inventory design review [Customer governance]	
MAS expectation	Review the inventory design to capture attributes for newer AI technologies and third-party AI.
FIOR support	The FI reviews its inventory design. FIOR extends the identity record as agent technologies evolve — for example, harness, permitted-model and operating-system attributes.
Evidence	Release notes

4.9 — Roles for inventorisation [Shared]	
MAS expectation	A designated control function with independent oversight of inventory accuracy and completeness.
FIOR support	Issuance and revocation of identities are restricted to authorised administrator roles, with multi-factor authentication enforced for privileged roles and every change recorded in the tamper-evident audit trail.
Evidence	Admin → Role matrix · Report → Audit

4.10 — Materiality methodology [Customer governance]	
MAS expectation	A consistent methodology for each AI use case, considering dependencies between AI systems and models.
FIOR support	The methodology is the FI's. FIOR supplies measurable inputs per agent: traffic volume, refusal rates and reasons, sensitive-data detections, tool use, models used and cost.
Evidence	Observe → Events · Report → Cost

4.11 — Inherent and residual risk [Customer governance]	
MAS expectation	Assess inherent and residual materiality; residual risk within appetite before deployment.
FIOR support	FIOR controls are mitigants the FI evaluates in its residual assessment. Shadow-mode policies record what a control would have done on live traffic before it is enforced, evidencing its effect.
Evidence	Control → Policies (shadow mode)

4.12 — Risk dimensions [Shared]	
MAS expectation	Minimally impact, complexity (including explainability and visibility into third-party AI) and reliance (including on third-party AI and the degree of autonomy).
FIOR support	Impact: sensitive-data and threat detections per agent. Complexity: provider and model recorded for each request. Reliance: autonomy evidenced by tool-call activity and human-approval requirements. Scoring remains the FI's.
Evidence	Observe → Events · Control → Approvals

© 2026 FIOR Group Limited
Visit our website: fior.group

4.13 — Roles for materiality assessment [Customer governance]	
MAS expectation	A designated control function to arbitrate or approve materiality, with documentation and regular review.
FIOR support	An FI governance assignment; FIOR records the resulting controls applied to each agent.
Evidence	Control → Policies

5 AI Life Cycle Controls

5.1 — Controls across the life cycle [Shared]	
MAS expectation	Robust controls across the full AI life cycle, with clear roles across business and functional units, reviewed for newer technologies.
FIOR support	FIOR governs the operational life cycle of each AI agent: enrolment and issuance, controlled operation under policy scoped to its operator, permitted models, domains and resources, expiry warnings and maximum certificate age, and revocation. Policies have their own life cycle: shadow mode, phased rollout, dual-control approval and versioned history. The model development life cycle is the FI's.
Evidence	Control → Agents · Control → Policies

5.2 — Proportionate controls and partial deployments [FIOR control]	
MAS expectation	Controls proportionate to materiality; mitigants such as limiting use where third-party disclosure is inadequate; governed pilots and phased rollouts with time and user limits and close monitoring.
FIOR support	Policies apply per agent, per path and per cohort, so use can be limited precisely. Phased rollouts can be confined to a percentage or named set of agents, restricted to time windows and run in shadow mode — directly supporting the time limits, user limits and monitoring MAS describes for partial deployments.
Evidence	Control → Policies

5.3 — Contingency plans and kill switches [Shared]	
MAS expectation	Contingency plans with fallback options for high-risk AI, tested under a range of conditions; clear, tested activation protocols for kill switches.
FIOR support	Revoking an agent's identity stops it at every gateway from the next policy synchronisation (configurable; 10 seconds in the FIOR Kubernetes appliance). Revocation lists published by the FI's own CA are checked every 15 seconds. Model routes let the FI direct a model to an alternative approved provider centrally, without changing agents. Gateways fail closed: they refuse traffic when their policy is out of date or, where configured, when they cannot write their audit record. Contingency plans and their testing are the FI's.
Evidence	Provider → Revocation · Control → Models · Observe → Events

© 2026 FIOR Group Limited
Visit our website: fior.group

5.4 — Data management [Shared]	
MAS expectation	Data fit for purpose, representative, high quality, classified, secure, privacy-compliant and traceable.
FIOR support	(e) Data security: inputs and outputs are protected in transit — credential detection, PII masking that fails closed, log redaction, a minimum TLS version and hybrid post-quantum key exchange to AI providers. (f) Privacy: personal data can be masked before it reaches an external AI service. (g) Each request's handling is recorded. Fitness, representativeness, quality and classification (a–d) are the FI's data governance.
Evidence	Control → Policies · Observe → Events

5.5 — Transparency and explainability controls [Shared]	
MAS expectation	Controls to ensure an appropriate level of transparency and explainability for each AI use case.
FIOR support	Every gateway decision is explainable: the record states the outcome, the reason and the policy that decided it, with that policy's version history. Explainability of model outputs is the FI's.
Evidence	Observe → Events · policy history

5.6 — Degree of transparency [Customer governance]	
MAS expectation	Transparency reflecting context and materiality: input features, key drivers, customer disclosure, consequences and redress, suited to the audience.
FIOR support	Customer disclosure and redress are FI processes. FIOR evidences where human approval was required and given for agent actions.
Evidence	Control → Approvals

5.7 — Defining fair outcomes [Customer governance]	
MAS expectation	Define fair outcomes and controls to identify and mitigate harmful bias.
FIOR support	Fairness definitions and assessments are the FI's. FIOR ensures the model assessed is the model in production: only approved models are permitted, deprecated models are refused, and system prompts can be pinned so the assessed configuration cannot change unnoticed.
Evidence	Control → Policies

5.8 — Fairness assessments [Customer governance]	
MAS expectation	Assess protected attributes, systematic disadvantage and document mitigations.
FIOR support	Performed by the FI on its model outcomes; FIOR's model and prompt pinning keeps assessments valid for the deployed configuration.
Evidence	Control → Policies

5.9(a) — Human oversight — roles [Shared]	
MAS expectation	Clear roles, escalation and decision-making for human oversight.
FIOR support	Approval requests are decided by authorised roles; responsible officers can be assigned and delegated review of high-impact decisions. The FI defines the roles.
Evidence	Control → Approvals · Corporate → Officer assignments

© 2026 FIOR Group Limited
Visit our website: fior.group

5.9(b) — Capabilities and authority to intervene [FIOR control]	
MAS expectation	Personnel equipped with the authority and ability to intervene.
FIOR support	Operators can revoke an agent, quarantine it across the fleet, require approval for designated tool calls and change policy under dual control.
Evidence	Provider → Revocation · Control → Policies

5.9(c) — Design for oversight [FIOR control]	
MAS expectation	Design AI to enable oversight, including mechanisms that escalate outputs for human review when pre-defined conditions are met.
FIOR support	This is how FIOR's approval control works: tool calls that meet conditions the FI defines — by tool, arguments or path — are held before the agent receives them and escalated for human approval. An approval is single-use, time-limited and bound to the exact agent, tool and arguments, so it cannot be reused for a different action.
Evidence	Control → Approvals

5.9(d) — Documentation and review of oversight [FIOR control]	
MAS expectation	Document logs and records from deployment, operation, monitoring and incident management; review oversight decisions and interventions, including near misses.
FIOR support	Gateway decisions, approvals, policy changes, revocations and officer sign-offs are recorded; the audit trail is hash-chained and verifiable; officer reviews export to CSV and PDF.
Evidence	Observe → Events · Report → Audit · Officer review export

5.10 — Third-party AI onboarding, testing and documentation [Shared]	
MAS expectation	Controls for third-party AI proportionate to materiality: contractual visibility of AI introduction and changes, testing in the FI's own context and data, compensatory testing, and documentation of suitability.
FIOR support	FIOR lets the FI evaluate a third-party AI service on its own live traffic under guardrails before full use: shadow-mode policies, phased cohorts and time windows, with every decision recorded as documentation of the testing. Contracts and the suitability decision are the FI's.
Evidence	Control → Policies · Observe → Events

5.11(a) — Third-party AI — transparency [Shared]	
MAS expectation	Assess third-party transparency, including independent certifications or external assessments; compensating controls such as additional testing or greater human oversight; limit, suspend or replace third-party AI where residual risk is outside appetite.
FIOR support	Compensating controls MAS lists — greater human oversight and controlled use — are enforced through approvals and policy, and a third-party AI service can be limited or suspended by policy, or replaced through model routes, at every gateway. For FIOR itself, the FI can rely on an independent external assessment rather than self-attestation (see Executive Summary).
Evidence	Control → Policies · Control → Models · Control → Approvals

© 2026 FIOR Group Limited
Visit our website: fior.group



5.11(b) — Supply chain assessments [Shared]	
MAS expectation	Supply chain risk assessment of third-party and open-source models, datasets and dependencies.
FIOR support	For third-party AI services, FIOR enforces approved models, pinned system prompts, pinned tool contracts and pinned upstream TLS certificates. For FIOR itself: a CycloneDX software bill of materials, dependency and container vulnerability scanning, secret scanning, and signed release artefacts.
Evidence	Control → Policies · FIOR release artefacts

5.11(c) — Concentration risk [Customer governance]	
MAS expectation	Assess over-reliance on key third-party AI providers, direct or indirect.
FIOR support	The FI's assessment. FIOR applies one identity and policy model across AI providers and self-hosted models, and model routes determine centrally which provider serves each model, so diversifying providers does not require a new control framework or changes to agents.
Evidence	Control → Models · Control → Policies

5.11(d) — Third-party AI change management [Shared]	
MAS expectation	Receive notifications of post-deployment updates or changes to third-party AI, and assess their impact.
FIOR support	Provider notifications are contractual. FIOR independently guards the FI's approved configuration against changes it can observe: requests for deprecated models are refused, and pinned system prompts, tool contracts and upstream certificates are refused if they change.
Evidence	Control → Policies

5.11(e) — Third-party contingency plans [Customer governance]	
MAS expectation	Contingency plans for third-party failure, unexpected behaviour or withdrawal of support.
FIOR support	FI plans; an alternative approved provider can be put in service centrally through model routes.
Evidence	Control → Models

5.11(f) — Legal agreements [Customer governance]	
MAS expectation	Performance guarantees, data protection, audit rights, notification when AI is introduced and updated.
FIOR support	FI contracts; FIOR's exportable records support the exercise of audit rights.
Evidence	Report → Audit

5.11(g) — Staff capabilities [Customer governance]	
MAS expectation	Awareness and capability of staff dealing with third-party AI.
FIOR support	FI training; FIOR provides product documentation and role-based guidance.
Evidence	FIOR documentation

© 2026 FIOR Group Limited
Visit our website: fior.group

5.11(h) — Complexity-based scrutiny [Customer governance]	
MAS expectation	More detailed assessment for complex third-party AI, e.g. AI agents.
FIOR support	FI assessment; tool-call governance and approvals give additional control over agentic AI.
Evidence	Control → Policies · Control → Approvals

5.12 — Selection of systems, models and features [Customer governance]	
MAS expectation	Document selection considerations commensurate with materiality, balancing performance against complexity, fairness and explainability.
FIOR support	FI selection process; FIOR enforces the outcome by permitting only the selected models.
Evidence	Control → Policies

5.13 — Newer and complex AI [Customer governance]	
MAS expectation	Weigh benefits against hallucination, opaqueness and security risks.
FIOR support	FI decision; FIOR's guardrails (5.15) mitigate the security risks of newer models.
Evidence	Control → Policies

5.14 — Evaluation and testing [Customer governance]	
MAS expectation	Evaluation measures and thresholds, testing approaches and overfitting mitigation before deployment.
FIOR support	Model evaluation is the FI's. FIOR records production behaviour — refusals, error responses, threat findings — that informs evaluation thresholds and monitoring.
Evidence	Observe → Events

5.15 — Guardrails and their effectiveness [FIOR control]	
MAS expectation	Controls and guardrails before deployment, which may include pilots and phased rollouts; test key Generative AI and agent failure modes and the effectiveness of guardrails.
FIOR support	Guardrails enforced inline: threat inspection of prompts, credential detection, PII masking, approved models and domains, pinned system prompts, prompt and output size limits, temperature limits, media-type restrictions and tool-call governance, with default deny. Shadow mode and phased cohorts measure each guardrail's effect on live traffic, and every decision is recorded as evidence of effectiveness.
Evidence	Control → Policies · Observe → Events

5.16(a) — Security [FIOR control]	
MAS expectation	Secured environments, hardened configuration, segmentation, input validation, API authentication, encryption and data loss prevention.
FIOR support	Hardened deployment: non-root containers, read-only root file systems, seccomp RuntimeDefault, all Linux capabilities dropped; an immutable-OS option. API authentication by agent identity; encryption by TLS with a policy-set minimum version and optional mutual TLS; input validation by request-size, prompt-length and media-type limits; data loss prevention by credential detection and PII masking.
Evidence	Deployment manifests · Control → Policies

© 2026 FIOR Group Limited
Visit our website: fior.group



5.16(b) — Access control [FIOR control]	
MAS expectation	Role-based access, multi-factor authentication, robust privileged access management, separation of duties.
FIOR support	Role-based access enforced by database row-level security; multi-factor authentication required for console access and enforced server-side for privileged administrator roles; dual control on policy changes; release of a quarantined agent restricted to super administrators. Provider credentials for AI services are held by FIOR rather than by the agents.
Evidence	Admin → Role matrix · Account → Security

5.16(c) — Third-party components [FIOR control]	
MAS expectation	Govern plugins and APIs, restrict data exposure, monitor for security and service issues.
FIOR support	Upstream AI services are restricted to approved hosts and regions with optional certificate pinning; tools and MCP servers are governed by allow and deny lists; outbound data is filtered; all calls are recorded.
Evidence	Control → Policies · Observe → Events

5.17 — Reproducibility and auditability [Shared]	
MAS expectation	Documentation, commensurate with materiality, sufficient for an independent party to understand and replicate the AI system and its testing.
FIOR support	FIOR records the operational configuration side: versioned policies with change notes, identity records, decision evidence under enforced retention, and optional signed provenance verifiable offline. Model development documentation is the FI's.
Evidence	Report → Audit · Observe → Provenance

5.18 — Pre-deployment review [Customer governance]	
MAS expectation	Review by parties not involved in development, proportionate to materiality.
FIOR support	FI review process. FIOR shadow mode and phased rollouts give reviewers evidence from live traffic, and dual control ensures a second person approves the policy change that puts an AI system into service.
Evidence	Control → Policies · Control → Approvals

5.19 — Independent validation (high risk) [Customer governance]	
MAS expectation	Formal independent validation of high-risk AI use cases before deployment.
FIOR support	FI validation function; FIOR evidence as for 5.18.
Evidence	Control → Policies

5.20 — Documented peer review (non-high risk) [Customer governance]	
MAS expectation	Peer review by qualified individuals not involved in development or deployment.
FIOR support	FI process; FIOR dual control evidences independent approval of the operating policy.
Evidence	Control → Approvals

© 2026 FIOR Group Limited
Visit our website: fior.group

5.21 — Review findings to approval body [Customer governance]	
MAS expectation	Report findings, limitations and conditions to the approval body and ensure they are actioned.
FIOR support	FI governance; conditions of use can be enforced as gateway policy.
Evidence	Control → Policies

5.22(a) — Secure design and deployment checklist [Shared]	
MAS expectation	Review architecture, segregation and least-privilege access; confirm critical controls before go-live.
FIOR support	FIOR identities are scoped to least privilege (permitted models, domains, resources). Gateways can be configured to refuse traffic if their audit record cannot be written, so logging is confirmed operational on every request. Go-live sign-off is the FI's.
Evidence	Control → Policies

5.22(b) — Testing and threat mitigation [Shared]	
MAS expectation	Vulnerability assessment, penetration testing and red teaming; safeguards such as input validation, throttling and anomaly detection.
FIOR support	FIOR provides the safeguards MAS names: input validation (size, length, media-type limits, threat inspection), throttling (rate limits, burst smoothing, concurrency, token budgets) and anomaly detection (cost anomalies, repeated requests, unusual path counts, error-rate thresholds). Security testing of the FI's AI systems is the FI's.
Evidence	Control → Policies

5.23(a) — Monitoring measures [FIOR control]	
MAS expectation	Metrics and tiered thresholds; drift checks; monitoring of reasoning processes, actions taken and tools used.
FIOR support	Every request and tool call is recorded with its decision, reason and deciding policy — covering the actions taken and tools used that MAS highlights for agentic workflows. Threshold checks enforce limits in real time; gateways warn when an agent's fallback policy is out of date. Model performance and drift metrics are the FI's.
Evidence	Observe → Events · Observe → Overview

5.23(b) — AI incident management and kill switches [FIOR control]	
MAS expectation	Report, track, escalate and resolve issues; kill switches or overrides for high-risk AI; user feedback channels.
FIOR support	Kill switch: revocation of an agent identity, applied at every gateway from the next policy synchronisation. Threat-triggered quarantine refuses an agent across every gateway sharing the deployment. Unregistered-AI alerts are raised for review. The incident process is the FI's.
Evidence	Provider → Revocation · Observe → Observations

© 2026 FIOR Group Limited
Visit our website: fior.group

5.23(c) — Roles and accountable person [Customer governance]	
MAS expectation	Assign monitoring and incident roles and an accountable person.
FIOR support	FI assignment, implemented through FIOR roles and responsible-officer assignment.
Evidence	Corporate → Officer assignments

5.23(d) — Documentation and access logging [FIOR control]	
MAS expectation	Records of monitoring and remediation; strictly controlled and logged access to AI systems and configuration; for Generative AI and agents, consider logging prompts, responses, model versions and reasoning.
FIOR support	Each request is recorded with the agent identity, the model requested, a preview of the prompt (with secrets masked where log redaction is enabled), token counts, the threat-inspection result, the tool calls made and the decision. Administrative actions are recorded in a server-side SHA-256 hash chain with immutability controls, in-console verification and full export. Gateway decision evidence cannot be deleted before its retention period. Optional provenance adds signed checkpoints and RFC 3161 timestamps verifiable offline.
Evidence	Observe → Events · Report → Audit · Observe → Provenance

5.23(e) — Training and awareness [Customer governance]	
MAS expectation	Train monitoring staff, and users to report unintended behaviour.
FIOR support	FI training programme.
Evidence	—

5.24 — Re-validation and review [Customer governance]	
MAS expectation	Re-validate or review AI as needed, proportionate to materiality; regular independent re-validation for high-risk AI; triggers include monitoring alerts and breaches.
FIOR support	FI review; FIOR's alerts, refusals and threshold breaches provide re-validation triggers, with a fleet-wide view of agents, decisions and cost and exportable evidence.
Evidence	Observe → Overview · Report → Cost · Report → Audit

5.25(a) — Scope of change [Shared]	
MAS expectation	A framework for the significance of changes; re-approval for significant changes; compensating controls such as enhanced monitoring where significant changes may occur without prior review.
FIOR support	The framework is the FI's. Dual control can be required for the policy changes the FI deems significant, and FIOR's pinning and per-request monitoring act as compensating controls where a third party changes its AI without prior review.
Evidence	Control → Policies · Observe → Events

© 2026 FIOR Group Limited
Visit our website: fior.group

5.25(b) — Change control [FIOR control]	
MAS expectation	Change control with appropriate human oversight; version control, traceability and rollback.
FIOR support	Policy changes are versioned with change notes and, under dual control, take effect only once a second authorised person approves. Model, system-prompt and tool-contract pinning detect and refuse unapproved changes to the AI configuration in use.
Evidence	Control → Policies (history) · Control → Approvals

5.25(c) — Dynamic updates [Shared]	
MAS expectation	Enhanced controls for AI that updates automatically.
FIOR support	Where a provider changes a model or tool definition, pinning and deprecated-model refusal keep the deployed configuration within what the FI approved.
Evidence	Control → Policies

5.26 — Retirement and decommissioning [Shared]	
MAS expectation	Controls for dependencies, data retention, secure removal and stakeholder notification.
FIOR support	Revocation or expiry removes an agent's access at every gateway; its identity history and decision evidence are retained. Data retention and stakeholder notification are the FI's.
Evidence	Provider → Revocation · Report → Audit

6 AI Capability & Capacity

6.1 — Personnel competence [Customer governance]	
MAS expectation	Competence, conduct, training and resourcing for personnel developing, deploying and using AI.
FIOR support	FI responsibility; FIOR provides product documentation and role-based guidance.
Evidence	FIOR documentation

6.2 — Capability programme review [Customer governance]	
MAS expectation	Regularly review capability-building programmes.
FIOR support	FI responsibility.
Evidence	—

6.3 — Technology infrastructure [Shared]	
MAS expectation	Adequate, resilient and secure infrastructure for AI.
FIOR support	FIOR deploys on hardened infrastructure for the governance plane — Kubernetes (K3s) including air-gapped installation, a dedicated appliance image, an immutable Talos-based OS, or an AWS machine image. The FI's AI compute infrastructure is outside FIOR's scope.
Evidence	Deployment options

© 2026 FIOR Group Limited
Visit our website: fior.group

6. Evidence artefacts

Reviewers will ask to be shown. The table below lists the evidence an FI running FIOR can produce, and the AIRG paragraphs each supports.

Artefact	Where it comes from	AIRG anchor
Tamper-evident audit trail (full-chain CSV) with chain verification	Report → Audit	4.3, 4.9, 5.9(d), 5.17, 5.23(d)
Gateway decision records: outcome, reason, identity, deciding policy	Observe → Events	1.10, 5.5, 5.15, 5.23(a), 5.23(d)
Policy version history with change notes	Control → Policies → history	3.5, 5.25(b)
Approval records for held tool calls and dual-control policy changes	Control → Approvals	2.5, 3.5, 5.9, 5.18
AI agent identity inventory and revocation records	Control → Agents · Provider → Revocation	4.5, 4.7, 5.3, 5.26
Unregistered-AI discovery and review decisions	Observe → Observations · Control → Sensors	1.10(f), 4.2, 4.4
Responsible-officer review export (CSV / PDF)	Corporate → Officer review	3.1, 3.5, 5.9(d)
Signed provenance evidence bundle, verifiable offline	fior-ctl provenance export / verify (provenance mode)	5.17, 5.23(d)
Cost and usage report	Report → Cost	3.2, 4.10, 5.24

7. Deployment options and rollout

Deployment	Best for
Kubernetes (K3s) bundle, including air-gapped installation	FIs with platform teams, and regulated network zones without internet access.
FIOR appliance image (ISO)	On-premises deployment where the gateway must run inside the FI's network.
Talos-based FIOR OS	Immutable-operating-system deployments with a minimal attack surface.
AWS machine image (bring your own licence)	AWS-first FIs running FIOR within their own AWS accounts.

Typical 12-week rollout

- **Weeks 1–2:** Deploy FIOR; establish the provider root CA or import the FI's own; enrol agents; deploy FIOR Sensors to discover AI use not yet enrolled.
- **Weeks 3–6:** Express the FI's AI policies — approved providers, models, domains and tools — as gateway policies in shadow mode; tune against live traffic.
- **Weeks 7–10:** Enforce in phased cohorts; enable dual control on policy changes, human approval for designated tool calls and evidence retention.

© 2026 FIOR Group Limited
Visit our website: fior.group

- **Weeks 11–12:** Full enforcement; export the audit trail, decision records and officer reviews as evidence for the FI's AIRG governance and review.

© 2026 FIOR Group Limited
Visit our website: fior.group

8. Appendix A — Glossary

Term	Meaning
AIRG	MAS Guidelines on Artificial Intelligence Risk Management, issued 7 October 2026 (consulted on as Consultation Paper P017-2025).
Agent identity	An X.509 certificate binding an AI agent to its provider, operator, permitted models and domains, with optional hybrid post-quantum signatures.
Provider	The organisation whose root CA issues agent identities — the FI itself or a business unit.
Operator	The team or entity running an AI agent that presents a FIOR identity.
Gateway	FIOR's policy decision point, which authorises every request against identity and policy.
Model route	The gateway setting that determines which AI provider or self-hosted model serves a given model; a request for a model with no route is refused.
Default deny	A request that no active policy permits is refused.
Shadow mode	A policy that records what it would have decided on live traffic without enforcing.
Shadow AI	Use of AI tools or services without the FI's knowledge or approval.
Dual control	A policy change takes effect only after approval by a second authorised person.
Revocation	Withdrawal of an agent identity, applied by every gateway at its next policy synchronisation.
Quarantine	Fleet-wide refusal of an agent for a set period after a threat detection.
PQC	Post-quantum cryptography: ML-DSA (FIPS 204) signatures and ML-KEM (FIPS 203) key exchange.

© 2026 FIOR Group Limited
Visit our website: fior.group

9. Appendix B — Gateway policy controls reference

Controls the FI can apply in a FIOR gateway policy, grouped by purpose.

Purpose	Controls
Identity	Model must match certificate · require post-quantum signature · require per-request signature · maximum certificate age · expiry warning · maximum clock skew · known agent types
Scope	Allowed models and domains · deprecated-model refusal · time windows · path scoping · required headers · phased cohorts · shadow mode
Content	Threat-inspection findings · credentials in prompt · PII masking · pinned system prompt · body, prompt and output limits · temperature limit · allowed media types · log redaction
Usage	Rate limit · burst smoothing · concurrency · model quotas · token budgets · cost anomaly · repeated requests · unusual path counts · error-response thresholds · stream limits
Upstream	Approved hosts and regions · certificate pinning · minimum TLS version
Tools	Allow and deny lists · argument restrictions · path restrictions · call quotas · loop detection · schema validation · pinned tool contracts · human approval
Platform	Refuse when audit cannot be written · evidence retention · fleet-wide threat quarantine · dual control · stale-policy warning

Tool-call controls apply to tool calls in Anthropic Messages and OpenAI Chat Completions responses. Where tool checks apply and a response cannot be inspected, the request is refused before it is sent.

© 2026 FIOR Group Limited
Visit our website: fior.group