Agentic AI Governance Framework Regulatory Mapping Tables

April 2026

Manage the Risks to Gain the Benefits

The Main AI Regulations And Standards Were Not Designed For Agentic AI

ISO 42001

Does not mention Agentic AI

- Defines an AI Management System in broad terms.
- Addresses AI's unique challenges, such as ethical considerations and transparency.

NIST AI RMF

Does not mention Agentic AI

- In either the core document AI RMF 1.0, 2023
- Or the Gen AI Profile AI-600-1, July 2024

Work is underway to address the gap

- March 2026, confirmation that monitoring agents must span functionality, operations, security, compliance, and human factors.
- AI Agent Standards Initiative launched.

EU AI Act – obligations without guidance

Does not mention Agentic AI

- Designed for static AI models, not agentic ones that evolve after deployment.
- Requires a risk management system but assumes a predictable risk profile.

AI agents fall in-scope by definition

- It's definition of an AI system is broad enough to capture agents, and the EU AI Office has confirmed this.
- The Act's prohibition of harmful manipulation applies to agent behaviour.
- The Act's 'high-risk' classification applies based on use case, not system type.

High-risk classification

- If you deploy an agent in one of the domains listed in Annex III — employment, credit, law enforcement, healthcare, border management, and others.
- It could be considered 'high-risk' and be subject to the full set of obligations under Chapter III.
- Yet the EU AI Act ***does not specify how*** deployers of agentic AI should comply at the operational level.

Compliance ≠ Governance

Agentic systems break four of the EU AI Act's core assumptions, which was written before agentic AI existed:

Oversight – Art. 14 requires human oversight, but as agent volumes grow, *it does not address* how to design for automation bias, alert fatigue, or *the degradation of oversight at scale*.

Risk management – Art. 9 requires a lifecycle risk system but *assumes a stable risk profile* that can be described at deployment. Yet, agentic systems learn and evolve after deployment.

Controls – the Act *does not distinguish between a policy* that describes a control and *a system that enforces it*. Documenting that an agent must not act beyond its boundaries is not the same as preventing it from doing so.

Cybersecurity – the Act does not mention prompt injection, multi-agent risks, and agent identity. Compliance with Art. 15 *will not close* these attack vectors.

Because of this, a firm that optimises for the Act's documentation requirements can simultaneously:

- ✓ Pass compliance checks.
- ✗ Be under regulatory scrutiny for inadequate agentic controls.

These are not criticisms of the Act

- They are known gaps in the regulatory landscape,
- That a governance framework designed for agentic AI will need to fill to:
  - Manage operational risk.
  - Avoid costly issues and negative attention if an agent goes wrong.

Closing The Gaps Requires 3 Shifts In How We Typically Design Governance

- Agents have their limitations, but their capacity and speed exceed humans.
- Because of this, effective governance will require not just specific controls but also 3 general shifts:

	From	To
1. Enforcement	**Descriptive governance** (e.g. "We have a policy that prohibits X").	**Operational governance**, (e.g. "The system cannot do X because of control Y.")
2. Continuity	**A project** with an end state.	**A permanent capability** – ongoing, evolving, and embedded into your operating model.
3. Scope	**Parallel governance** for humans vs AI.	**Unified governance** for 'combined operations' that comprise human and non-human workers.

Conclusion for risk managers

For many firms, achieving this will mean:

- Evolving their governance and controls
- For 'combined operations' of humans and non-humans.

Agentic AI Governance Framework

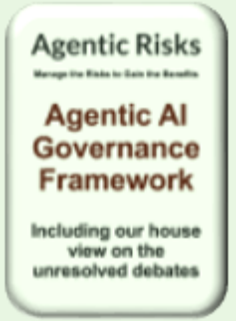

Download your copy here

Evolve Your Governance and Controls
Everything You Need In One Place

5 Foundations of Traditional AI Governance That Remain Relevant

1. Policy and Principles.
2. AI Inventory and Lifecycle Management.
3. Data Governance.
4. Human Oversight and Controls.
5. Accountability, Evidence, and Audit Trails.

9 New Components You Will Need To Govern Agentic AI

1. Agent Identity, Permissions, and Prohibitions.
2. Dynamic Risk Assessment.
3. Pre-Execution Boundaries, Control, and Security.
4. Reasoning Chain Integrity.
5. Multi-Agent Governance.
6. Lifecycle Governance for Systems That Learn and Adapt.
7. Human Oversight Redesigned for Automation Bias.
8. End-User Responsibility and Transparency.
9. Workforce Integration and Organisational Readiness.

Additional Features

1. Security embedded at every step.
2. 'Controls checklists':
   - Specific and implementable.
   - Software agnostic.
   - Map to the Enterprise-Wide Agentic AI Risk Controls.
3. Mapping tables: EU AI Act, ISO, and NIST.
4. 'House views' in every section, because understanding theory has more value if you can also implement it in practice.
5. Practical mitigations for the key unresolved industry debates.

ISO/IEC 42001:2023 : Agentic Risks' Agentic AI Governance Framework

Summary Mapping Table

Regulatory Requirement	Ref	Agentic Risks	Sub-Heading
Organisational Context	Cl. 4	✅ Covered	1.1 Policy and Principles
Leadership and AI Policy	Cl. 5	✅ Covered	1.1 Policy and Principles
Risk Assessment and Treatment	Cl. 6.1	✅ Covered	2.2 Dynamic Risk Assessment
AI Objectives and Planning	Cl. 6.2	✅ Covered	1.1 Policy and Principles
Support (Competence, Awareness, Resources)	Cl. 7	✅ Covered	2.9 Organisational Readiness for Combined Operations
Operational Planning and Control	Cl. 8.1	✅ Covered	2.3 Pre-Execution Boundaries, Control, and Security
AI Impact Assessment	Cl. 8.2	✅ Covered	2.2 Dynamic Risk Assessment
AI System Lifecycle	Cl. 8.3	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Data for AI Systems	Cl. 8.4	✅ Covered	1.3 Data Governance
Monitoring and Performance Evaluation	Cl. 9.1	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Internal Audit	Cl. 9.2	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Management Review	Cl. 9.3	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Improvement and Corrective Action	Cl. 10	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
AI Policies	A.2	✅ Covered	1.1 Policy and Principles
Internal Organisation and Roles	A.3	✅ Covered	1.1 Policy and Principles
Resources for AI Systems	A.4	✅ Covered	1.2 AI Inventory and Lifecycle Management
AI System Lifecycle Controls	A.5	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Data Governance for AI	A.6	✅ Covered	1.3 Data Governance
Information for Interested Parties	A.7	✅ Covered	2.8 End-User Responsibility and Transparency
Use of AI Systems	A.8	✅ Covered	2.7 Human Oversight Redesigned for Automation Bias
Third-Party AI Relationships	A.9	✅ Covered	3.1 Liability Across the Agentic Value Chain

Source: ISO/IEC 42001:2023 — Artificial Intelligence — Management System. International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC), December 2023.

NIST AI RMF 1.0 : Agentic Risks' Agentic AI Governance Framework

Summary Mapping Table

Regulatory Requirement	Ref	Agentic Risks	Sub-Heading
Legal and Regulatory Compliance	GV-1.1	✅ Covered	1.1 Policy and Principles
Trustworthy AI Integration	GV-1.2	✅ Covered	1.1 Policy and Principles
Risk-Based Decision Making	GV-1.3	✅ Covered	2.2 Dynamic Risk Assessment
Transparent Risk Management	GV-1.4	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Organisational Risk Tolerance	GV-1.5	✅ Covered	1.1 Policy and Principles
Policies for AI Risks	GV-1.6	✅ Covered	1.1 Policy and Principles
Processes for AI Risks	GV-1.7	✅ Covered	2.2 Dynamic Risk Assessment
Roles and Responsibilities	GV-2.1	✅ Covered	1.1 Policy and Principles
Human-AI Configuration Policies	GV-2.2	✅ Covered	2.7 Human Oversight Redesigned for Automation Bias
Feedback Mechanisms	GV-3.1	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Diversity and Inclusion in Risk	GV-3.2	✅ Covered	2.2 Dynamic Risk Assessment
Safety-First Organisational Culture	GV-4.1	✅ Covered	2.9 Organisational Readiness for Combined Operations
Documenting AI Risks and Impacts	GV-4.2	✅ Covered	2.2 Dynamic Risk Assessment
Organisational Risk Policies	GV-5.1	✅ Covered	1.1 Policy and Principles
Incident Response Policies	GV-5.2	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Policies for AI Risk Management	GV-6.1	✅ Covered	3.1 Liability Across the Agentic Value Chain
AI Risk Management Practices	GV-6.2	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Context Establishment	MP-1.1	✅ Covered	1.2 AI Inventory and Lifecycle Management
Organisational Risk Tolerance	MP-1.5	✅ Covered	2.2 Dynamic Risk Assessment
Scientific and Technical Knowledge	MP-2.3	✅ Covered	2.2 Dynamic Risk Assessment
AI Risk Identification	MP-3.1	✅ Covered	2.2 Dynamic Risk Assessment
Risks to Third Parties	MP-4.1	✅ Covered	2.8 End-User Responsibility and Transparency
Performance Metrics	MS-1.1	✅ Covered	2.2 Dynamic Risk Assessment
Testing and Evaluation	MS-2.5	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Bias and Fairness Testing	MS-2.6	✅ Covered	1.3 Data Governance
Explainability and Interpretability	MS-3.3	✅ Covered	2.4 Reasoning Chain Integrity
Deployment Monitoring	MS-4.1	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Prioritised Risk Treatment	MG-1.3	✅ Covered	2.2 Dynamic Risk Assessment
Incident Response Plans	MG-2.2	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
AI Risk Remediation	MG-3.1	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Continual Improvement	MG-4.1	✅ Covered	3.5 Agentic Governance is an Ongoing Activity

Source: NIST AI 100-1 — Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology, U.S. Department of Commerce, January 2023.

EU AI Act High-Risk Obligations : Agentic Risks' Agentic AI Governance Framework

Summary Mapping Table

Ref	Regulatory Requirement	Agentic Risks	Sub-Heading
Art. 6	Classification rules for high-risk AI systems	✅ Covered	1.1 Policy and Principles
Art. 8	Compliance with requirements	✅ Covered	1.1 Policy and Principles
Art. 9	Risk management system	✅ Covered	2.2 Dynamic Risk Assessment
Art. 10	Data and data governance	✅ Covered	1.3 Data Governance
Art. 11	Technical documentation	✅ Covered	2.2 Dynamic Risk Assessment
Art. 12	Record-keeping / automatic logging	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Art. 13	Transparency and information to deployers	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Art. 14	Human oversight	✅ Covered	1.4 Human Oversight and Controls
Art. 15	Accuracy, robustness, and cybersecurity	✅ Covered	2.3 Pre-Execution Boundaries, Control, & Security
Art. 16	Obligations of providers of high-risk AI systems	✅ Covered	General
Art. 17	Quality management system	✅ Covered	General
Art. 18	Documentation keeping	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Art. 19	Automatically generated logs	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Art. 20	Corrective actions and duty to inform authorities of serious incidents	✅ Covered	2.6 Lifecycle Governance for Systems That Learn & Adapt
Art. 21	Cooperation with competent authorities	✅ Covered	1.5 Accountability, Evidence, and Audit Trails
Art. 25	Responsibilities along the AI value chain	✅ Covered	3.1 Liability Across the Agentic Value Chain
Art. 26	Obligations of deployers of high-risk AI systems	✅ Covered	General
Art. 27	Fundamental rights impact assessment (FRIA)	✅ Covered	2.2 Dynamic Risk Assessment
Arts. 43, 47-49	Conformity assessment, EU declaration of conformity, CE marking, and EU database registration	✅ Covered	General
Art. 72	Post-market monitoring plan	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt
Art. 73	Reporting of serious incidents	✅ Covered	

Source: Regulation (EU) 2024/1689, Official Journal 13 June 2024 Chapter III – High-Risk AI System Obligations vs Agentic Risks' Agentic AI Governance Framework.
Exclusions: Articles 7 and 28–39 of Chapter III are excluded from this mapping because they impose obligations exclusively on the European Commission (Art. 7, which governs amendments to the high-risk use case list) and on Member State governments and notified bodies (Arts. 28–39, which govern the designation and oversight of conformity assessment bodies) — neither creates any obligation for organisations deploying AI agents.

Operational Governance Requirements

The Rules Are Not Yet Complete

Operational Requirement	Why It Is Vital	Agentic Risks	EU AI Act	ISO 42001	NIST AI RMF
Agent Identity, Permissions, and Prohibitions.	No identity means no accountability.	✅ Covered (2.1)	Silent	Silent	Silent
Pre-Execution Boundaries, Control, and Security.	Stop bad actions before they happen.	✅ Covered (2.3)	Silent	Silent	Silent
Reasoning Chain Integrity.	Correct outputs can obscure flawed reasoning, which is a risk signal.	✅ Covered (2.4)	Silent	Silent	Silent
Multi-Agent Governance.	Errors travel fast in connected systems.	✅ Covered (2.5)	Silent	Silent	Silent
Lifecycle Governance for Systems That Learn and Adapt.	Escalating costs signal runaway agents. Inputs change. Behaviour drifts. Governance must follow.	✅ Covered (2.6)	Silent	Silent	Silent
Human Oversight Redesigned for Automation Bias.	Without realistic design, human oversight can degrade.	✅ Covered (2.7)	Silent	Silent	Silent
Organisational Readiness for Combined Operations.	You cannot introduce non-human workers without preparing the human workers.	✅ Covered (2.9)	Silent	Silent	Silent
Liability Across the Agentic Value Chain.	You need to be able to allocate liability when an agentic system causes harm.	✅ Covered (3.1)	Silent	Silent	Silent

Appendix: Detailed Mapping Tables

ISO/IEC 42001:2023 : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
Organisational Context	Cl. 4	Organisations must identify internal/external factors affecting AI governance, define AIMS scope, and understand stakeholder expectations regarding responsible AI use.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** requires board-level policy defining risk tiers and accountability, directly addressing organisational context and stakeholder obligations.
Leadership and AI Policy	Cl. 5	Top management must demonstrate commitment, establish an AI policy, assign accountability roles, and integrate AI governance into organisational strategy.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** calls for board-level policy with named executives, risk tiers, and governance arrangements, mirroring Cl. 5 leadership requirements.
Risk Assessment and Treatment	Cl. 6.1	Organisations must identify, assess, and treat AI-related risks, defining risk criteria and acceptable risk levels before and after deployment.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment's** pre- and post-deployment risk assessments map directly to Cl. 6.1, including risk classification, scoring, and treatment planning.
AI Objectives and Planning	Cl. 6.2	Organisations must set measurable AI governance objectives, plan how to achieve them, and document the planning process.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** establishes governance principles and risk appetite recommends formal, measurable objectives and documented planning and monitoring to achieve them.
Support (Competence, Awareness, Resources)	Cl. 7	Organisations must ensure adequate competence, awareness programmes, documented information, and resources to operate the AIMS effectively.	✅ Covered	2.9 Organisational Readiness for Combined Operations	**2.9 Organisational Readiness for Combined Operations** addresses workforce training, de-skilling, and readiness broadly and recommends formal competence documentation, awareness programmes, and evidenced resource allocation.
Operational Planning and Control	Cl. 8.1	Organisations must plan, implement, and control processes needed to meet AIMS requirements and implement risk treatment plans.	✅ Covered	2.3 Pre-Execution Boundaries, Control, and Security	**2.3 Pre-Execution Boundaries, Control, and Security** requires machine-enforceable execution boundaries, pre-deployment controls, and operational governance plans – directly satisfying Cl. 8.1.
AI Impact Assessment	Cl. 8.2	Organisations must assess the impact of AI systems on individuals, groups, and society, particularly for high-risk deployments.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment's** framework includes fundamental rights impact assessment (FRIA) requirements for high-risk systems, meeting Cl. 8.2.
AI System Lifecycle	Cl. 8.3	Organisations must manage AI systems across their full lifecycle – design, development, deployment, monitoring, and decommissioning – with documented controls.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** dedicates an entire component to lifecycle governance including gradual rollout, continuous monitoring, re-testing, and version control.

Source: ISO/IEC 42001:2023 — Artificial Intelligence — Management System. International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC), December 2023.

ISO/IEC 42001:2023 : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
Data for AI Systems	Cl. 8.4	Organisations must govern the quality, provenance, privacy compliance, and bias characteristics of data used in AI systems.	✅ Covered	1.3 Data Governance	**1.3 Data Governance's Data Governance** section explicitly covers data quality, provenance, GDPR, bias monitoring – directly satisfying Cl. 8.4.
Monitoring and Performance Evaluation	Cl. 9.1	Organisations must establish processes to monitor AIMS performance, measure effectiveness against objectives, and report results to management.	✅ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence, and Audit Trails** requires KRIs, anomaly thresholds, audit trails, and management reporting, meeting Cl. 9.1's performance evaluation requirements.
Internal Audit	Cl. 9.2	Organisations must conduct periodic, impartial internal audits to verify AIMS conformity with ISO 42001 requirements and organisational policies.	✅ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence, and Audit Trails** references regular audits of oversight quality and policy compliance and explains the importance of an impartial, scheduled internal audit programme.
Management Review	Cl. 9.3	Top management must periodically review the AIMS to assess its continuing suitability, adequacy, and effectiveness.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** calls for continuous monitoring, regular re-testing, and a formal management review process.
Improvement and Corrective Action	Cl. 10	Organisations must address AIMS nonconformities, take corrective action, and drive continual improvement of the management system.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt's** continuous lifecycle governance, incident management, and post-incident learning directly correspond to Cl. 10 improvement obligations.
AI Policies	A.2	Establish and maintain a formal, board-approved AI policy providing management direction for responsible AI development and use.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** mandates a board-level AI policy with named accountabilities, risk tiers, and governance arrangements – meeting A.2 in full.
Internal Organisation and Roles	A.3	Define and assign AI governance roles, responsibilities, and reporting channels; establish mechanisms for raising concerns.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** requires named executive accountabilities, oversight roles, and governance arrangements that map to A.3 control objectives.
Resources for AI Systems	A.4	Document all resources required for AI systems – data, tooling, computing infrastructure, and human expertise – across the lifecycle.	✅ Covered	1.2 AI Inventory and Lifecycle Management	**1.2 AI Inventory and Lifecycle Management** recommends structured documentation of all resources - data, tooling, compute, skills - at each lifecycle stage.
AI System Lifecycle Controls	A.5	Apply documented controls at each lifecycle stage: design, development, testing, deployment, monitoring, and decommissioning.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** articulates five governance moments (pre-deployment, gradual rollout, continuous monitoring, re-testing, management review), addressing A.5 lifecycle controls.

Source: ISO/IEC 42001:2023 — Artificial Intelligence — Management System. International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC), December 2023.

ISO/IEC 42001:2023 : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
Data Governance for AI	A.6	Govern data quality, provenance, privacy, consent, and bias across all data used in AI system development and operation.	✅ Covered	1.3 Data Governance	**1.3 Data Governance** explicitly covers all A.6 data governance dimensions including quality, provenance, GDPR, bias monitoring, and amplification risks.
Information for Interested Parties	A.7	Provide stakeholders with appropriate, transparent information about AI systems – capabilities, limitations, and decision-making processes.	✅ Covered	2.8 End-User Responsibility and Transparency	**2.8 End-User Responsibility and Transparency** calls for both internal and external users map directly to A.7 requirements for stakeholder information.
Use of AI Systems	A.8	Define and control how AI systems are used, including approved use cases, user responsibilities, and override mechanisms.	✅ Covered	2.7 Human Oversight Redesigned for Automation Bias	**2.7 Human Oversight Redesigned for Automation Bias** addresses use through calibrated oversight requirements, approved action types, and meaningful human intervention points per A.8.
Third-Party AI Relationships	A.9	Manage risks from suppliers, partners, and customers involved in the AI lifecycle, including contractual and due-diligence obligations.	✅ Covered	3.1 Liability Across the Agentic Value Chain	**3.1 Liability Across the Agentic Value Chain** addresses vendor contractual requirements in the liability section and the importance of systematic third-party AI governance.

Source: ISO/IEC 42001:2023 — Artificial Intelligence — Management System. International Organization for Standardization (ISO) / International Electrotechnical Commission (IEC), December 2023.

NIST AI RMF 1.0 : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
Legal and Regulatory Compliance	GV-1.1	Understand, document, and actively manage AI-related legal and regulatory requirements throughout the AI system lifecycle.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** explicitly references regulatory obligations, requiring organisations to identify and manage legal requirements.
Trustworthy AI Integration	GV-1.2	Integrate trustworthy AI characteristics – safety, fairness, explainability, accountability – into organisational policies and practices.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles'** governance principles embed fairness, explainability, and accountability requirements across all five foundations and nine new components.
Risk-Based Decision Making	GV-1.3	Establish procedures to determine appropriate levels of risk management activity based on organisational risk tolerance.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** requires risk tiers tied to autonomy, action-space, and reversibility, with controls calibrated accordingly – satisfying GV-1.3.
Transparent Risk Management	GV-1.4	Govern risk management processes through transparent, documented policies, procedures, and mechanisms.	✅ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence, and Audit Trails** requires defensible, reconstructable evidence of AI decisions and transparent governance arrangements, meeting GV-1.4.
Organisational Risk Tolerance	GV-1.5	Document and communicate organisational risk tolerance for AI, informing decisions on acceptable AI system deployment and use.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** calls for board-level risk appetite statements covering each category of agentic risk, directly satisfying GV-1.5.
Policies for AI Risks	GV-1.6	Implement policies addressing risks across the AI lifecycle, including development, deployment, and decommissioning.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** requires policies covering permitted autonomy levels, in/out-of-scope workflows, and governance for all lifecycle stages.
Processes for AI Risks	GV-1.7	Establish processes for identifying, assessing, and managing AI risks across the organisation and AI system lifecycle.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment's** pre- and post-deployment assessments, continuous monitoring, and KRI thresholds meet GV-1.7 process requirements.
Roles and Responsibilities	GV-2.1	Document clear roles, responsibilities, and communication lines for AI risk management throughout the organisation.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** calls for named executive owners, risk management roles, and governance arrangements, satisfying GV-2.1.
Human-AI Configuration Policies	GV-2.2	Define roles and responsibilities for human oversight within AI-enabled systems, including human-AI task allocation.	✅ Covered	2.7 Human Oversight Redesigned for Automation Bias	**2.7 Human Oversight Redesigned for Automation Bias** focuses on defining human roles in agentic workflows and redesigning human oversight for automation bias.
Feedback Mechanisms	GV-3.1	Establish processes to capture and act on feedback about AI system performance from internal and external sources.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** covers continuous monitoring, anomaly detection, and a stakeholder feedback mechanism.

Source: NIST AI 100-1 — Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology, U.S. Department of Commerce, January 2023.

NIST AI RMF 1.0 : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
Diversity and Inclusion in Risk	GV-3.2	Prioritise workforce diversity, equity, inclusion, and accessibility in AI risk management processes.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** recommends a variety of inputs and perspectives.
Safety-First Organisational Culture	GV-4.1	Foster a critical thinking and safety-first mindset across the organisation for AI design, deployment, and use.	✅ Covered	2.9 Organisational Readiness for Combined Operations	**2.9 Organisational Readiness for Combined Operations** stresses the need for engaged, informed human oversight and workforce integration.
Documenting AI Risks and Impacts	GV-4.2	Require teams to document AI risks and impacts, and communicate them across the organisation.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** mandates structured risk documentation across pre- and post-deployment assessments, including scoring, controls, and KRI design.
Organisational Risk Policies	GV-5.1	Develop and implement organisational risk policies to guide AI risk management practices.	✅ Covered	1.1 Policy and Principles	**1.1 Policy and Principles** requires comprehensive policies covering risk tiers, autonomy levels, and governance arrangements – directly satisfying GV-5.1.
Incident Response Policies	GV-5.2	Establish policies and procedures for responding to AI-related incidents, including escalation, communication, and resolution.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** covers incident classification, kill switches, escalation, and post-incident learning within its lifecycle governance component.
Policies for AI Risk Management	GV-6.1	Organisations understand, manage, and document AI risk management requirements across the supply chain.	✅ Covered	3.1 Liability Across the Agentic Value Chain	**3.1 Liability Across the Agentic Value Chain** addresses vendor contracting and liability, and calls for a systematic supply-chain AI risk management policy.
AI Risk Management Practices	GV-6.2	AI risk and benefit information is shared across the organisation and with relevant AI actors.	✅ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence, and Audit Trails** requires audit trails, management reporting, and external disclosures, enabling internal AI risk information sharing.
Context Establishment	MP-1.1	Identify and document the intended purpose, scope, and operational context of each AI system, including stakeholder impacts.	✅ Covered	1.2 AI Inventory and Lifecycle Management	**1.2 AI Inventory and Lifecycle Management** requires an AI inventory mapped to risk tier and purpose, which satisfies MP-1.1 context establishment requirements.
Organisational Risk Tolerance	MP-1.5	Risks and benefits of AI system deployment are understood and documented based on organisational risk tolerance.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment's** methodology examines action-space, reversibility, and blast radius relative to defined risk appetite – satisfying MP-1.5.
Scientific and Technical Knowledge	MP-2.3	Scientific and technical knowledge about AI is used to inform risk identification and assessment.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** draws on current agentic AI knowledge and calls for the integration of evolving technical research into risk assessments.

Source: NIST AI 100-1 — Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology, U.S. Department of Commerce, January 2023.

NIST AI RMF 1.0 : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
AI Risk Identification	MP-3.1	AI system risks are identified, mapped across the lifecycle, and linked to potential impacts on individuals and society.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment's** risk assessment maps risks across five categories with lifecycle-stage controls, directly satisfying MP-3.1 requirements.
Risks to Third Parties	MP-4.1	AI risks and impacts to third parties, including end-users, are evaluated and addressed.	✅ Covered	2.8 End-User Responsibility and Transparency	**2.8 End-User Responsibility and Transparency** addresses third-party transparency, disclosure obligations, and the impact of agents on end-users – satisfying MP-4.1.
Performance Metrics	MS-1.1	Metrics and methods are established to measure AI system performance and risk-related characteristics.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** requires KRIs, anomaly thresholds, and behavioural baselines for each deployed agent, meeting MS-1.1's requirements.
Testing and Evaluation	MS-2.5	Testing, evaluation, verification, and validation (TEVV) activities are conducted throughout the AI lifecycle.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** calls for pre-deployment testing, gradual rollout validation, continuous monitoring, and regular re-testing.
Bias and Fairness Testing	MS-2.6	AI systems are tested for bias, fairness, and potential for discriminatory outcomes.	✅ Covered	1.3 Data Governance	**1.3 Data Governance** calls for bias monitoring and fairness testing, meeting MS-2.6.
Explainability and Interpretability	MS-3.3	AI system explainability and interpretability are evaluated relative to the intended use and risk level.	✅ Covered	2.4 Reasoning Chain Integrity	**2.4 Reasoning Chain Integrity** addresses inspectable artefacts and the limits of explainability.
Deployment Monitoring	MS-4.1	Risks are monitored and tracked in deployed AI systems, with documented processes for updating risk assessments.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt's** continuous monitoring, automated anomaly detection, and drift thresholds directly satisfy MS-4.1 deployment monitoring requirements.
Prioritised Risk Treatment	MG-1.3	Responses to identified AI risks are prioritised based on impact and likelihood, with documented risk treatments.	✅ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** requires structured risk scoring, control design, and KRI implementation, providing a documented risk treatment process.
Incident Response Plans	MG-2.2	Incident response plans and procedures are in place, tested, and updated for AI-related events.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** requires defined incident classification, escalation protocols, kill switches, and post-incident review.
AI Risk Remediation	MG-3.1	Risks and impacts are managed through documented remediation activities, with monitoring of outcomes.	✅ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance for Systems That Learn and Adapt** calls for execution boundaries, anomaly detection, rollback triggers, and post-incident learning.
Continual Improvement	MG-4.1	Risk management activities are continuously improved based on incident learnings, testing results, and evolving risks.	✅ Covered	3.5 Agentic Governance is an Ongoing Activity	**3.5 Agentic Governance is an Ongoing Activity's** explicitly addresses agentic governance as a permanent, evolving capability – satisfying MG-4.1.

Source: NIST AI 100-1 — Artificial Intelligence Risk Management Framework (AI RMF 1.0). National Institute of Standards and Technology, U.S. Department of Commerce, January 2023.

EU AI Act High-Risk Obligations : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Obligation	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
1. Classification of AI systems as high-risk	Art. 6	Classification rules for high-risk AI systems	AI systems listed in Annex III (biometrics, critical infrastructure, education, employment, essential services, law enforcement, migration, justice) are high-risk. Classification is by use case and context, not architecture.	☑ Covered	1.1 Policy and Principles	**1.1 Policy & Principles** adopts the EU AI Act's four-tier risk structure and requires board-level classification of each agent by risk tier.
				☑ Covered	2.2 Dynamic Risk Assessment	**1.2 Dynamic Risk Assessment** explicitly warns that use-case classification alone is insufficient for agentic systems and requires assessment of autonomy level, action-space, and reversibility alongside use case.
2. Requirements for high-risk AI systems (Arts. 8–15)	Art. 8	Compliance with requirements	High-risk AI systems must comply with the requirements in Section 2 throughout their lifecycle. Providers bear primary responsibility; deployers share responsibility for their use.	☑ Covered	1.1 Policy and Principles	**1.1 Policy & Principles** establishes that governance is a continuous, lifecycle-wide obligation.
	Art. 9	Risk management system	Providers must establish, implement, document, and maintain a continuous risk management system across the full lifecycle - covering risk identification, estimation, evaluation, and mitigation. Testing required before market placement.	☑ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** directly addresses continuous, lifecycle-wide agentic risk management. Explicitly rejects static, one-time assessment. Component 6 (Lifecycle Governance) reinforces continuous testing, with gradual rollout, automated anomaly detection, and regular re-testing as distinct governance moments.
				☑ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance** reinforces continuous testing, with gradual rollout, automated anomaly detection, and regular re-testing as distinct governance moments.
	Art. 10	Data and data governance	Training, validation, and test datasets must be subject to governance covering relevance, representativeness, completeness, and freedom from errors. Bias examination required. Data minimisation and privacy compliance mandatory.	☑ Covered	1.3 Data Governance	**1.3 Data Governance** covers data quality, provenance, lineage, privacy compliance (GDPR / CCPA), and bias monitoring. The Framework notes that agentic systems intensify these requirements due to broader data access and active exfiltration risk.

Source: Regulation (EU) 2024/1689, Official Journal 13 June 2024 Chapter III – High-Risk AI System Obligations vs Agentic Risks' Agentic AI Governance Framework.
Exclusions: Articles 7 and 28–39 of Chapter III are excluded from this mapping because they impose obligations exclusively on the European Commission (Art. 7, which governs amendments to the high-risk use case list) and on Member State governments and notified bodies (Arts. 28–39, which govern the designation and oversight of conformity assessment bodies) — neither creates any obligation for organisations deploying AI agents.

EU AI Act High-Risk Obligations : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Obligation	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
2. Requirements for high-risk AI systems (Arts. 8–15)	Art. 11	Technical documentation	Pre-deployment technical documentation package required, covering system description, architecture, components, tool integrations, training methodology, test results, and monitoring measures. Must be kept current per Annex IV.	☑ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** integrates the Art. 11 technical documentation requirement, specifying that high-risk agents require formal pre-deployment documentation covering intended purpose, architecture, components, tool integrations, protocols, training / validation / testing methodology, monitoring measures, and test results - maintained throughout operational life.
	Art. 12	Record-keeping / automatic logging	High-risk AI systems must technically allow automatic event logging throughout their operational life, sufficient to reconstruct the period of use, inputs, outputs, and human review events.	☑ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence & Audit Trails** requires reconstructable evidence of AI decisions including inputs, model version, human review record, and outcome.
				☑ Covered	1.5 Accountability, Evidence, and Audit Trails	**2.4 Reasoning Chain Integrity** requires logging at each reasoning stage before progression to the next. Both 1.5 and 2.4 directly address the Art. 12 logging obligation.
	Art. 13	Transparency and provision of information to deployers	High-risk AI systems must be designed to be sufficiently transparent for deployers to understand and appropriately use outputs. Providers must supply instructions covering intended purpose, accuracy metrics, known limitations, human oversight measures, and maintenance requirements.	☑ Covered	2.8 End-User Responsibility and Transparency	**2.8 End-User Responsibility & Transparency** requires clear disclosure of agent capabilities, limitations, and escalation paths.
				☑ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence, and Audit Trails** requires defensible, reconstructable evidence. The Framework's treatment goes beyond the Act's transparency obligation for those firms who want to cover the specific challenges of agentic reasoning opacity, including latent reasoning and the limits of explainability.

Source: Regulation (EU) 2024/1689, Official Journal 13 June 2024 Chapter III – High-Risk AI System Obligations vs Agentic Risks' Agentic AI Governance Framework.
Exclusions: Articles 7 and 28–39 of Chapter III are excluded from this mapping because they impose obligations exclusively on the European Commission (Art. 7, which governs amendments to the high-risk use case list) and on Member State governments and notified bodies (Arts. 28–39, which govern the designation and oversight of conformity assessment bodies) — neither creates any obligation for organisations deploying AI agents.

EU AI Act High-Risk Obligations : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Obligation	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
2. Requirements for high-risk AI systems (Arts. 8–15)	Art. 14	Human oversight	High-risk AI systems must be designed so humans can effectively oversee them during use. Oversight measures must be commensurate with risk level and autonomy. Humans must be able to understand, monitor, interpret, override, and halt the system. Awareness of automation bias required.	☑ Covered	1.4 Human Oversight and Controls	**1.4 Human Oversight & Controls** addresses the Act's requirements for human oversight. It also notes that monitoring ≠ control, so agentic governance also requires the ability to refuse or halt activity, which the Framework covers in Section 2.3
				☑ Covered	2.7 Human Oversight Redesigned for Automation Bias	**2.7 Human Oversight Redesigned for Automation Bias** provides implementable guidance for mitigating the degradation of oversight quality over time, alert fatigue, reviewer capacity limits, and the structural design of oversight to counteract automation bias.
	Art. 15	Accuracy, robustness, and cybersecurity	High-risk AI systems must achieve appropriate accuracy, be resilient to errors and faults, include redundancy solutions, and resist unauthorised third-party manipulation. Systems that continue to learn must mitigate feedback loop risks.	☑ Covered	2.3 Pre-Execution Boundaries, Control, and Security	**2.3 Execution Boundaries & Pre-Action Control** covers prompt injection and input sanitisation as specific cybersecurity controls.
				☑ Covered	2.3 Pre-Execution Boundaries, Control, and Security	**2.5 Multi-Agent Governance** addresses cascading failures and orchestrator subversion as robustness concerns.
				☑ Covered	2.3 Pre-Execution Boundaries, Control, and Security	**2.6 Lifecycle Governance** feedback loop covers the risks in continuously learning systems and drift monitoring.

Source: Regulation (EU) 2024/1689, Official Journal 13 June 2024 Chapter III – High-Risk AI System Obligations vs Agentic Risks' Agentic AI Governance Framework.
Exclusions: Articles 7 and 28–39 of Chapter III are excluded from this mapping because they impose obligations exclusively on the European Commission (Art. 7, which governs amendments to the high-risk use case list) and on Member State governments and notified bodies (Arts. 28–39, which govern the designation and oversight of conformity assessment bodies) — neither creates any obligation for organisations deploying AI agents.

EU AI Act High-Risk Obligations : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Obligation	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
3. Obligations of providers and deployers (Arts. 16–27)	Art. 16	Obligations of providers of high-risk AI systems	Providers must: comply with Section 2 requirements; establish a quality management system; keep technical documentation; retain automatically generated logs; take corrective action and inform authorities of serious incidents; cooperate with authorities; affix CE marking; register in EU database.	☑ Covered	General	**The Framework** covers the substantive governance obligations (risk management, documentation, logging, oversight, incident reporting). Formal market-placement obligations - CE marking (Art. 48), EU declaration of conformity (Art. 47), and EU database registration (Art. 49) - are included for completeness but noted as being prerequisites for lawful deployment, rather than governance activities.
	Art. 17	Quality management system	Providers must implement a documented QMS covering: regulatory compliance strategies; design / development techniques; testing and validation procedures; technical standards; data management; risk management; post-market monitoring; accountability procedures; resource management.	☑ Covered	General	**The Framework** embeds the substantive Art. 17 QMS elements into its sections on Policy and Principles, Dynamic Risk Assessment, Lifecycle Governance, and Accountability, and explains that 'together, they address the required content needed to satisfy Art. 17.'
	Art. 18	Documentation keeping	Providers must keep technical documentation and quality management system documentation available to authorities for 10 years after the system is placed on the market.	☑ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence & Audit Trails** states that the EU AI Act 'requires a 10-year retention period and specific documentation types' as part of the regulatory scrutiny framing.
	Art. 19	Automatically generated logs	Providers must retain automatically generated logs to the extent they are under their control. Deployers must retain logs for at least 6 months.	☑ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence & Audit Trails** and **2.4 Reasoning Chain Integrity** require logging of agent actions, tool calls, decision points, and reasoning stages.
				☑ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance** adds prompt change logs and version control. The 6-month minimum deployer retention period is noted in the EU AI Act note under 'Obligations of deployers.'

Source: Regulation (EU) 2024/1689, Official Journal 13 June 2024 Chapter III – High-Risk AI System Obligations vs Agentic Risks' Agentic AI Governance Framework.
Exclusions: Articles 7 and 28–39 of Chapter III are excluded from this mapping because they impose obligations exclusively on the European Commission (Art. 7, which governs amendments to the high-risk use case list) and on Member State governments and notified bodies (Arts. 28–39, which govern the designation and oversight of conformity assessment bodies) — neither creates any obligation for organisations deploying AI agents.

EU AI Act High-Risk Obligations : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Obligation	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
3. Obligations of providers and deployers (Arts. 16–27)	Art. 20	Corrective actions and duty to inform authorities of serious incidents	Providers must immediately take corrective action and inform national authorities when a high-risk AI system presents a risk or causes a serious incident. Deployers must inform the provider and market surveillance authorities without undue delay.	☑ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance** includes a defined incident classification protocol covering the Act's specific external reporting obligations. The protocol covers when a malfunction constitutes a serious incident under the Act, the requirement to notify the provider and relevant national market surveillance authority, and the threshold (risk to health, safety, or fundamental rights under Art. 79(1)).
	Art. 21	Cooperation with competent authorities	Providers and deployers must cooperate with national competent authorities and provide all information and documentation necessary to demonstrate compliance.	☑ Covered	1.5 Accountability, Evidence, and Audit Trails	**1.5 Accountability, Evidence & Audit Trails** directly supports Art. 21 by stating that governance should produce defensible, reconstructable evidence for regulatory scrutiny. The framework explicitly positions regulatory defensibility as a primary governance objective, not a secondary one.
	Art. 25	Responsibilities along the AI value chain	Importers, distributors, and other parties who modify or place high-risk AI systems on the market take on provider obligations. Contracts must reflect these allocations.	☑ Covered	3.1 Liability Across the Agentic Value Chain	**3.1 Liability Across the Agentic Value Chain** addresses the absence of settled liability allocation and recommends explicit documentation of accountability chains and vendor contracts that address liability for agent behaviour specifically. Aligns with Art. 25's chain-of-responsibility principle.
	Art. 26	Obligations of deployers of high-risk AI systems	Deployers must: use systems per provider instructions; assign qualified human oversight; ensure input data relevance; monitor operation; inform providers and authorities of incidents; retain logs for ≥6 months; inform workers; verify EU database registration before use; cooperate with authorities.	☑ Covered	General	**The Framework** covers the substantive deployer obligations (human oversight, monitoring, data quality, incident reporting), explicitly acknowledging three remaining deployer-specific requirements that sit outside the governance framework: (1) EU database registration verification before first use, (2) worker notification obligation, and (3) 6-month log retention minimum. It characterises these as legal prerequisites rather than governance activities.
	Art. 27	Fundamental rights impact assessment (FRIA)	Public bodies, financial institutions, and insurers must conduct a FRIA before first use of a high-risk AI system, covering: deployment processes; period and frequency of use; affected persons; specific risks of harm; human oversight measures; risk mitigation measures. Must be updated on material changes.	☑ Covered	2.2 Dynamic Risk Assessment	**2.2 Dynamic Risk Assessment** integrates the FRIA requirement directly, specifying that deployers that are public bodies, financial institutions, or insurers subject to the EU AI Act must conduct a FRIA before first use of any high-risk AI system.

EU AI Act High-Risk Obligations : Agentic Risks' Agentic AI Governance Framework

Section	Ref	Obligation	Regulatory Requirement	Agentic Risks	Sub-Heading	Justification
4. Notifying Authorities and Notified Bodies (Arts. 28–39)		Out of scope: covers how notifying authorities must be organised at the member state level.				
5. Standards, Conformity Assessment, Certificates, Registration (Arts. 40–49)	Arts. 43 / 47 / 48 / 49	Conformity assessment, EU declaration of conformity, CE marking, and EU database registration	Providers of Annex III high-risk systems must conduct a conformity assessment before placement on market (Art. 43), draw up an EU declaration of conformity (Art. 47), affix CE marking (Art. 48), and register in the EU database (Art. 49). Deployers must verify registration before first use.	☑ Covered	General	**The Framework** explicitly lists CE marking, conformity assessment, and declaration of conformity under 'market-placement obligations' as prerequisites for lawful deployment. These are correctly characterised as legal prerequisites that are not governance activities but are noted for completeness.
	Art. 72	Post-market monitoring plan	Providers must establish and document a post-market monitoring plan before placement on market and actively gather and review data on system performance throughout operational life.	☑ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance** - the Framework's continuous monitoring requirement (gradual rollout, automated anomaly detection, drift thresholds, regular re-testing, and prompt / knowledge source version control) directly corresponds to and exceeds the Art. 72 post-market monitoring obligation for agentic systems.
	Art. 73	Reporting of serious incidents	Providers must report serious incidents to the market surveillance authority of the Member State where the incident occurred, without undue delay and within defined timeframes.	☑ Covered	2.6 Lifecycle Governance for Systems That Learn and Adapt	**2.6 Lifecycle Governance** - the Framework's incident classification protocol covers the Art. 73 external reporting obligation, including the threshold criterion (risk to health, safety, or fundamental rights under Art. 79(1)) and the requirement to notify the relevant national market surveillance authority without undue delay.

Source: Regulation (EU) 2024/1689, Official Journal 13 June 2024 Chapter III – High-Risk AI System Obligations vs Agentic Risks' Agentic AI Governance Framework.
Exclusions: Articles 7 and 28–39 of Chapter III are excluded from this mapping because they impose obligations exclusively on the European Commission (Art. 7, which governs amendments to the high-risk use case list) and on Member State governments and notified bodies (Arts. 28–39, which govern the designation and oversight of conformity assessment bodies) — neither creates any obligation for organisations deploying AI agents.