LEAMAN CRELLIN

ISITC EUROPE CIC

AI Governance and Oversight Guidance for Financial Services Firms

Principles, Guidance and Examples

Authors: Ian Perham, ISITC Europe CIC and Katharine Leaman, Leaman Crellin

March 2026 - Version 1

This document is supplied for the benefit of financial services firms operating in the United Kingdom. It draws on input from ISITC Europe CIC, The Investing and Saving Alliance (TISA) and the FCA Mills Review process.

This paper has been written as practical market guidance and does not constitute legal advice. Firms should take their own legal and regulatory advice on matters specific to their circumstances.

Foreword

Artificial intelligence is no longer a technology of the future. It is already embedded in core processes across financial services: in fraud detection, credit decisioning, transaction monitoring, customer communication and increasingly in agentic systems that act autonomously on behalf of firms and consumers alike.

The pace of adoption has outrun the development of formal regulatory requirements in the United Kingdom. The FCA is expected to publish comprehensive guidance on AI audit trails and governance following the Mills Review in Summer 2026. The UK AI Bill, expected in May 2026, is anticipated to introduce high-risk categorisation standards broadly aligned with the EU AI Act. Financial services firms cannot wait for that guidance before putting governance frameworks in place.

This document provides a set of Principles for AI governance in financial services. It is modelled on recognised market practice frameworks such as the Global Foreign Exchange Code and draws on the FCA's established approach to Operational Resilience and the Senior Managers and Certification Regime. The aim is to provide firms with practical, proportionate guidance now.

The Principles are structured to be accessible to firms of all sizes and authorisation types. Where requirements differ between dual-regulated firms and FCA solo-regulated firms, this is noted explicitly. The guidance reflects the current regulatory environment and draws on industry discussions including ISITC Europe roundtables on real-time compliance and TISA's response to the Mills Review.

Governance should enable growth. The Principles presented here are not intended to slow down AI adoption. They are intended to provide the confidence that allows firms to move faster, with appropriate accountability built in from the start.

Contents

Section 1 Introduction and Scope

1.1 Purpose

This document sets out Principles for the governance and oversight of artificial intelligence systems deployed by financial services firms in the United Kingdom. It addresses firms regulated by the Financial Conduct Authority (FCA), including solo-regulated firms, and firms that are also subject to Prudential Regulation Authority (PRA) oversight.

The Principles are intended to be applied proportionately. A firm deploying AI only for internal document summarisation faces different risks from a firm using AI agents to make real-time credit decisions. The guidance distinguishes between these contexts throughout.

1.2 Scope

These Principles apply to the governance and oversight of any AI system deployed by a firm, including:

- generative AI tools used for content creation, advice drafting, or customer communication.
- machine learning models used for fraud detection, credit scoring, transaction monitoring, or market surveillance.
- agentic AI systems that take autonomous actions on behalf of the firm or its customers
- AI systems provided by third parties and integrated into the firm's operations.
- AI tools adopted by staff without formal IT approval, sometimes referred to as shadow AI.

These Principles do not cover AI used purely for personal productivity purposes with no interaction with firm data, customer data, or regulated processes.

1.3 How to Use This Document

Each Principle is structured as follows:

- a statement of the Principle, which represents the standard firms should meet.
- guidance, which explains what the Principle means in practice and how firms can demonstrate compliance.
- examples, which illustrate application in specific contexts.
- notes on firm type, where requirements differ between dual-regulated and FCA solo-regulated firms.

The examples are illustrative. They are drawn from scenarios discussed across the financial services industry and are not exhaustive. Firms should exercise judgement in applying the Principles to their own circumstances.

Section 2 The Regulatory Context

2.1 The Current UK Position

The UK currently operates a principles-based, outcomes-focused approach to AI regulation. No single AI-specific statute applies to financial services firms at the time of publication. However, a range of existing regulatory obligations are directly relevant to how AI must be governed:

- Consumer Duty requires firms to deliver good outcomes for retail customers. AI systems used in retail contexts must be consistent with this obligation regardless of whether AI-specific rules have been published.
- the Senior Managers and Certification Regime (SMCR) requires that responsibility for material risks is allocated to named Senior Manager Function holders. AI risks are material risks for most firms deploying them at scale.
- Operational Resilience rules require firms to identify, test and protect important business services. AI systems embedded in important business services fall within this framework.
- the Critical Third Parties (CTP) regime, now active, enables the FCA and PRA to oversee the resilience of key technology providers. Firms must assess whether AI infrastructure providers qualify as CTPs or as third parties requiring equivalent oversight.
- the FCA's transaction reporting requirements, addressed in MarketWatch 81, directly engage data quality and governance standards that underpin AI-fed systems.

The FCA published the Mills Review call for input in late 2025 and is expected to publish a formal response in Summer 2026. That response is anticipated to include the first comprehensive FCA guidance on AI audit trails. Firms should treat this document as a foundation on which FCA guidance will build, not as a substitute for it.

2.2 The EU AI Act

The EU AI Act came into force in August 2024. Its provisions are being phased in over a period running to August 2027. The Act is horizontal in scope: it applies across all sectors of the economy. Financial services is not treated as a distinct category within the Act's structure, though many AI applications in financial services will fall within the Act's high-risk classification.

The EU AI Act is relevant to UK firms for two reasons. First, any firm with customers in the European Economic Area, or using EU-based service providers, must assess whether the Act applies to its AI systems. Second, the UK AI Bill, expected in May 2026, is anticipated to adopt high-risk categorisation standards broadly aligned with the Act's approach. The Act therefore provides a useful framework for classification now.

There is an important distinction between the EU AI Act and this guidance. The Act is a horizontal instrument applicable across all sectors. This guidance is vertical: it is directed specifically at financial services firms and addresses the particular risks, regulatory obligations, and governance structures relevant to that sector. The Act sets minimum standards. This guidance goes further where the nature of financial services requires it.

2.3 Key Milestones

March 2026: UK Government AI and Copyright Report published. This clarifies how firms may use AI models trained on protected data.

May 2026: UK AI Bill expected. Anticipated to introduce statutory high-risk AI categorisation broadly aligned with the EU AI Act.

Summer 2026: FCA formal response to the Mills Review expected. Likely to include the first comprehensive FCA guidance on AI audit trails.

August 2026: EU AI Act prohibition and general-purpose AI obligations fully in force.

August 2027: EU AI Act high-risk system obligations fully in force for systems placed on the market before August 2026.

Section 3 Risk Classification Framework

These Principles start with risk. A firm cannot build appropriate governance structures without first understanding what risks its AI systems present. The classification framework below provides a starting point. It draws on the EU AI Act high-risk categorisation and adapts it for the specific context of UK financial services.

Firms should treat this framework as a minimum. They should apply their own risk assessment processes and seek appropriate expertise where the classification of a system is unclear.

Risk Level	Category	Example Use Cases	Minimum Controls Required
High	Consumer-facing autonomous decision	AI agent making real-time investment recommendations; AI system determining credit eligibility; automated advice suitability assessment; AI-driven pricing in insurance	Board approval; named SMF accountability; independent model validation; human override capability; Consumer Duty impact assessment; audit trail; regulatory notification where required
High	Financial crime and surveillance	AI transaction monitoring for AML purposes; market surveillance systems; fraud detection systems with automated block or flag capability	Named SMF accountability; independent validation; documented false positive rate management; MLRO oversight; full audit trail
High	Operational AI in important business services	AI embedded in settlement processing; AI managing client reporting obligations; AI in regulatory transaction reporting	Operational resilience mapping; business continuity planning; third-party CTP assessment; kill switch capability; SMF24 (dual-regulated) or equivalent oversight
Medium	Internal decisioning support	AI drafting credit memos for human review; AI summarising regulatory intelligence; AI generating compliance reports for human sign-off	Documented human review process; accuracy monitoring; named business owner; periodic validation
Medium	Customer-facing non-advisory	AI-driven chatbot for general enquiries; AI routing customer complaints; AI personalising marketing communications	Consumer Duty review; bias testing; escalation to human agents; monitoring for vulnerability signals
Low	Internal productivity	AI summarising meeting notes; AI drafting internal communications; AI assisting with code review in non-production environments	Acceptable use policy; data classification controls preventing input of confidential or customer data; periodic review

Firms should note that classification is not static. An AI system that begins as low risk can become high risk if its scope expands or if it becomes integrated into a process that is subsequently classified as an important business service. Risk classifications should be reviewed at least annually and following any material change to an AI system's function or deployment context.

Section 4 The Principles

The eight Principles below represent the standards that firms should meet in governing and overseeing AI systems. They are intended to be read as a whole. A firm that addresses only some of the Principles will not have an adequate governance framework.

Principle 1 Risk Identification and Classification

Firms should identify, classify, and document all AI systems in use. The level of governance applied to each system should be proportionate to the risks it presents.

GUIDANCE

1.1 Firms should maintain a firm-wide inventory of AI systems. The inventory should record, for each system, its purpose, the data it processes, the decisions it influences or makes, and whether it involves customer data or affects regulated activities. The inventory should include systems procured from third parties and systems adopted by staff without formal IT approval.

1.2 Each system in the inventory should be assigned a risk classification consistent with the framework in Section 3. The classification should reflect the impact of a failure or error, the degree of autonomy the system exercises, the reversibility of decisions it makes, and the extent of consumer or market exposure.

1.3 High-risk systems should not be deployed without prior sign-off from the relevant Senior Manager Function holder (see Principle 3) and, where applicable, Board approval. Sign-off should be recorded and maintained as part of the AI governance record.

1.4 The inventory and classifications should be reviewed at least annually and following any material change to a system's function, data inputs, operating context, or regulatory environment.

Example 1.1 Shadow AI Inventory Exercise

A UK asset manager conducts a firm-wide review of AI adoption. The review discovers that three front-office analysts have been using a commercial generative AI tool to draft client communication, inputting client portfolio information. The tool has not been reviewed by IT, Legal or Compliance.

The firm classifies this as medium risk given the customer data involved and the potential for Consumer Duty implications. The tool is suspended pending a formal risk assessment. Following assessment, it is approved for use under a restricted acceptable use policy that prohibits input of identified client data.

The inventory is updated and the COO's function assumes oversight responsibility.

Principle 2 Board Accountability and AI Strategy

The Board bears ultimate responsibility for the firm's approach to AI and should set the firm's AI strategy and risk appetite. The Board should receive regular reporting on AI risk exposure and be equipped to provide informed challenge.

GUIDANCE

2.1 The Board should approve a written AI strategy that sets out the firm's objectives for AI adoption, the risk appetite within which AI will be deployed, and the governance structures that will apply. The strategy should be reviewed at least annually.

2.2 AI risks should be integrated into the firm's enterprise risk management framework. They should appear in the firm's risk register with defined risk owners, tolerances, and escalation triggers.

2.3 The Board should receive a regular AI governance report. This report should include an AI heatmap showing the firm's current AI use cases by risk classification, material incidents or near-misses involving AI systems in the reporting period, changes to the AI inventory or risk classifications, and an assessment of the firm's AI risk exposure relative to its risk appetite.

2.4 Board members do not need to be AI technical experts. They do need to understand the material risks the firm's AI systems present and to be able to provide informed challenge to the executives responsible. Boards should consider whether they have access to appropriate expertise, whether through board membership, a Board advisory resource, or through the firm's second line functions.

2.5 The Board's approach to AI governance should be consistent with its approach to Operational Resilience and third-party risk. In each case the Board sets the appetite, receives the reporting, and holds senior management accountable for delivery.

Example 2.1 Board AI Heatmap

A retail bank presents its Board with a quarterly AI heatmap. The heatmap shows twelve current AI use cases arranged by risk classification. Three are classified as high risk: an AI-driven mortgage affordability model, an AI transaction monitoring system and an AI system used to route vulnerable customer contacts.

The heatmap shows that the mortgage affordability model has had two material errors in the quarter, both corrected before customer impact. The Board instructs the CRO to commission an independent model validation review and requests a report at the next meeting.

The Board approves the deployment of a new medium-risk AI tool for complaint summarisation, subject to a Consumer Duty assessment being completed before go-live.

Analogy: The FCA's Operational Resilience rules established that the Board owns important business service decisions and sets impact tolerances. AI governance should follow the same principle. The Board does not manage AI systems. It owns the framework within which they operate.

Note: Firms subject to the Senior Managers Regime should ensure that Board-level AI accountability is reflected in the Chair's Statement of Responsibilities where relevant.

Principle 3 Senior Management Accountability under SMCR

Firms should clearly allocate AI accountability to named Senior Manager Function holders. Accountability should follow the business function that commissions AI use. One designated Senior Manager should own the technical review, challenge, and due diligence function for all AI systems across the firm.

GUIDANCE

3.1 The Senior Managers and Certification Regime requires that responsibility for material risks is allocated to named Senior Manager Function holders. AI systems that are used in, or affect, regulated activities are material risks for most firms deploying them at scale. Firms should treat AI accountability in the same way as they treat accountability for Operational Resilience: the commissioning business area owns the outcome; a designated technical SMF owns the integrity of the system.

3.2 Accountability for the business outcomes of an AI system rests with the Senior Manager Function holder responsible for the function that commissions its use. An AI system used in the credit function is the responsibility of the SMF responsible for credit. An AI system used in financial crime monitoring is the responsibility of the SMF responsible for financial crime. This reflects how accountability flows in Operational Resilience and is consistent with FCA expectations of a single point of accountability.

3.3 One designated Senior Manager should own the firm's technical due diligence function for AI. This function includes the review, validation, and challenge of AI systems before deployment and on an ongoing basis. In practice this is likely to be the holder of the COO or technology-focused SMF function, as they will typically hold the expertise or access to expertise required. In dual-regulated firms this is likely to be SMF24. In FCA solo-regulated firms this responsibility should be allocated to the most senior operational or technology executive and recorded in their Statement of Responsibilities.

3.4 Where the business SMF commissioning an AI system does not have the technical expertise to assess the system independently, they should be able to rely on the technical review of the designated technical SMF or on external specialist expertise engaged for the purpose. That reliance should be documented. The commissioning SMF remains accountable for the outcome.

3.5 Statements of Responsibilities should be updated to reflect AI accountability. The FCA expects to see AI oversight addressed in SoRs. A system capable of causing a systemic outage or a Consumer Duty breach requires explicit accountability.

3.6 The firm's AI governance committee should be chaired or co-chaired by the business-accountability SMF and the technical-accountability SMF. The committee should review all new AI deployments before they move from proof of concept to live operation.

Important note on firm type: SMF4 (Chief Risk Officer) and SMF24 (Chief Operations Officer) are PRA-designated functions. They do not exist in FCA solo-regulated firms. Firms authorised by the FCA only should allocate the risk function and operational function accountability to the relevant executive director (SMF3) or CEO (SMF1) and record this explicitly in their Statements of Responsibilities.

Note: The ISITC Europe 2026 analysis identified SMF24 as the natural home for operational integrity and technology oversight of AI in dual-regulated firms, with SMF4 owning model risk and SMF16 owning consumer outcomes. This structure is adopted in this guidance.

SMF	Function	AI Accountability	Firm Type
SMF1	Chief Executive	Sets ethical tone and AI strategy; accountable for firm's overall AI risk posture; responsible for regulatory relationship on AI matters	All firms
SMF3	Executive Director	In FCA solo firms: may hold risk and operational accountability for AI where no SMF4 or SMF24 exists; must be specified in SoR	FCA solo firms
SMF4	Chief Risk Officer	Owns model risk management; commissions independent model validation; ensures AI operates within risk appetite; oversees model bias and accuracy monitoring	Dual-regulated firms only (PRA designation)
SMF16	Compliance Oversight	Signs off regulatory compliance of AI use cases; owns Consumer Duty assessment for AI; commissions human oversight of AI decisions in regulated contexts; monitors FCA	All firms

SMF	Function	AI Accountability	Firm Type
		perimeter for AI-delivered regulated activities	
SMF17	Money Laundering Reporting Officer	Accountable for AI systems used in financial crime detection and monitoring; approves false positive management policies; owns AML audit trail for AI-generated alerts	All firms with AML obligations
SMF24	Chief Operations Officer	Designated technical SMF; owns due diligence, validation, and challenge function for all AI systems; responsible for operational resilience of AI; ensures kill switch capability; oversees third-party AI provider resilience; updates SoR to reflect AI operational integrity	Dual-regulated firms only (PRA designation)
Commissioning SMF	Varies by use case	Accountable for business outcomes of AI deployed in their function; signs off risk classification; relies on SMF24 (or equivalent) for technical assessment; cannot delegate ultimate accountability	All firms

Example 3.1 Accountability Mapping for an AI Suitability Tool

A wealth management firm introduces an AI system to assist advisers in generating initial suitability assessments for review. The system uses client data, risk profile information, and product data to produce a draft recommendation which the adviser reviews and either approves, amends or rejects.

The commissioning SMF is the SMF responsible for the advisory function (in this case an SMF3 executive director). They are accountable for whether the tool produces suitable outcomes and whether advisers are properly trained to challenge it.

The technical accountability sits with the COO (an SMF3 in this FCA solo firm with COO responsibilities noted in their SoR). They commission an independent model validation, document the data inputs and commission a Consumer Duty impact assessment before deployment.

SMF16 (Compliance Oversight) signs off the regulatory framework within which the tool operates and sets the standards for human review of AI-generated assessments.

The Board approves deployment of this high-risk AI system and receives quarterly reporting on override rates and outcome quality.

Principle 4 DUE Diligence and Model Risk Management

Before any AI system is deployed in a live environment, it should be subject to robust due diligence. Due diligence for high-risk systems should include independent technical validation. Firms should not rely solely on the assessment provided by the system's developer or supplier.

GUIDANCE

4.1 Firms should establish a due diligence standard appropriate to the risk classification of each AI system. At a minimum, due diligence should assess: the data on which the model was trained and whether that data is representative of the population to which it will be applied; the model's accuracy, bias characteristics and error rates; the conditions under which the model may produce unreliable or harmful outputs; the system's behaviour when data quality is poor or incomplete; and the controls available to the firm to identify and correct errors.

4.2 For high-risk AI systems, validation should be conducted by a function or individual that is independent of those who developed or procured the system. Where the firm does not have in-house technical expertise of sufficient depth to conduct this validation, it should engage external specialist expertise. The reliance on external expertise should be documented and the commissioning SMF retains accountability for the outcome.

4.3 AI systems procured from third parties require equivalent due diligence to those developed in-house. A firm cannot transfer its regulatory accountability to a vendor. Vendor documentation and conformity assessments required under the EU AI Act should be obtained and reviewed but do not substitute for the firm's own assessment.

4.4 Firms should maintain a model risk management framework that encompasses AI systems. The framework should define how models are approved, how they are monitored in production, how material changes are governed, and how models are decommissioned. This framework should be owned by the designated technical SMF.

4.5 The output of the due diligence process should be documented and retained. Documentation should be sufficient to allow the firm to explain to the FCA or to a customer, why a particular AI system was deployed, what assessment was conducted, and what controls are in place.

Example 4.1 Due Diligence for a Third-Party AI Financial Crime System

A UK bank procures an AI-driven transaction monitoring system from a specialist fintech vendor. The system is classified as high risk.

The bank's technical SMF (SMF24) commissions a due diligence exercise. The vendor provides technical documentation including training data provenance, model architecture, false positive and false negative rates against industry benchmark datasets, and their EU AI Act conformity assessment.

The bank's model risk team reviews the documentation and identifies that the training data underrepresents certain transaction patterns common in the bank's customer base. They require the vendor to conduct supplementary testing on a representative dataset before deployment.

The bank also commissions an external technology consultant to review the vendor's resilience arrangements and confirm that the system can be operated by the bank independently if the vendor relationship fails.

SMF17 (MLRO) reviews the framework for human review of AI-generated alerts and approves the false positive management policy. Deployment is approved by the AI governance committee and noted in the Board report.

Principle 5 Data Governance, INTEGRITY, and Audit Trails

AI systems should be underpinned by robust data governance. Firms should understand the data journeys that feed their AI systems, ensure data quality and integrity, and maintain audit trails that allow AI-influenced decisions to be reconstructed and explained.

GUIDANCE

5.1 Firms should map the data journeys that supply each AI system. This means understanding, for each system, what data is ingested, where that data originates, how it is processed before reaching the AI system, and how the AI system's outputs are used in downstream processes. The ISITC Europe roundtables on real-time compliance identified that traceability across business functions is foundational to identifying and addressing AI-related data quality issues.

5.2 The FCA's MarketWatch 81 publication identified five root causes of poor data quality in transaction reporting: change management, reporting process and logic design, data governance, control framework, and governance, oversight, and resourcing. These root causes apply equally to AI systems that feed on regulatory data. Firms should assess their AI data governance against each of these dimensions.

5.3 Firms must be able to explain AI-influenced decisions. This is not merely a regulatory expectation: it is a prerequisite for meaningful human oversight and for the Consumer Duty obligation to support customer understanding. Firms should ensure that the outputs of AI systems are logged with sufficient contextual information to allow the decision to be reconstructed, the data inputs identified, and the model version recorded.

5.4 For AI systems used in regulated activities, audit trail requirements should align with the record-keeping obligations that apply to the equivalent human-conducted activity. An AI system conducting functions equivalent to investment advice must generate and retain records consistent with MiFID II record-keeping requirements.

5.5 Data quality controls should be preventative as well as detective. Upstream controls that prevent poor-quality data entering AI systems are more effective than controls that identify errors after AI-driven decisions have been made. This is particularly important in near-real-time processing environments where the settlement cycle leaves limited time for error correction.

5.6 Firms should implement data standards for AI inputs. The BCBS 239 framework for risk data aggregation provides a useful reference for the principles that should underpin AI data governance: accuracy, completeness, timeliness, and adaptability.

Example 5.1 Audit Trail Requirements for AI in Transaction Monitoring

A firm's AI transaction monitoring system flags a transaction as suspicious. The alert is reviewed by a human analyst who decides not to file a Suspicious Activity Report.

The audit trail must record: the transaction data that was processed, the version of the AI model that generated the alert, the alert score and the factors that contributed to it, the identity of the human analyst who reviewed it, the basis for the decision not to file, and the date and time of each step.

If the FCA subsequently investigates the transaction, the firm must be able to produce this record and demonstrate that the human review was genuine and substantive, not merely a rubber stamp of the AI output.

Note: The FCA has signalled in its Mills Review consultation that it expects AI audit trails in financial crime contexts to meet this standard. Firms should prepare now.

Principle 6 Operational Resilience and Third-Party Risk

AI systems embedded in important business services should be subject to the same operational resilience standards as other critical systems. Firms should assess whether AI

infrastructure providers require oversight under the Critical Third Parties regime or equivalent standards.

GUIDANCE

6.1 Firms should assess whether any of their AI systems are embedded in, or necessary for the continuity of, an important business service as defined under the FCA's and PRA's Operational Resilience rules. Where they are, the AI system must be within scope of the firm's resilience framework: mapped, tested, and subject to impact tolerance assessment.

6.2 Every high-risk AI system should have a documented kill switch: a mechanism to suspend or revert the AI's decision-making function and replace it with a manual or legacy process. The kill switch must be operable by a named individual with appropriate authority and must be tested as part of the firm's operational resilience testing programme.

6.3 Firms should assess their AI infrastructure providers against the criteria for the Critical Third Parties regime. Providers of foundational model infrastructure (including large language model APIs), specialist AI platforms and AI-enabled data services may qualify as, or behave like, critical third parties regardless of their formal designation. Firms should apply equivalent oversight standards even where a provider is not formally designated.

6.4 Firms should monitor and manage concentration risk in AI infrastructure. Dependency on a single large language model provider or AI platform creates risks that are analogous to IT outsourcing concentration risk. This risk should be assessed and reported to the Board.

6.5 The designated technical SMF should maintain oversight of vendor resilience for all high-risk AI systems. This includes reviewing vendor business continuity plans, exit provisions in contracts, and the firm's ability to operate its processes without the vendor during an outage.

Example 6.1 Kill Switch and Manual Override

A consumer lender uses an AI system to make initial credit decisions for online applications. The system processes approximately 2,000 applications per day and operates with minimal human intervention.

The firm's resilience framework requires that if the AI system fails or produces results outside defined parameters, the firm can revert to a manual review process within four hours. A team of credit analysts is trained and kept available to perform manual assessments.

The kill switch is tested quarterly. The most recent test identified that the manual process took six hours rather than four. The firm required the vendor to improve system recovery time and internally cross-trained additional analysts.

SMF24 owns the kill switch procedure and confirms its operability to the Board annually.

Principle 7 Consumer Outcomes and Market Integrity

Firms deploying AI in consumer-facing processes must ensure that AI-driven outcomes are consistent with Consumer Duty and do not cause foreseeable harm. Firms should actively monitor the regulatory perimeter to ensure AI does not deliver regulated activities outside the regulatory framework.

GUIDANCE

7.1 Consumer Duty requires firms to act to deliver good outcomes across the four outcome areas: products and services, price and value, consumer understanding, and consumer support. AI systems deployed in consumer-facing processes must be assessed against each of these outcome areas before deployment and monitored against them in production.

7.2 Firms should give specific attention to consumer vulnerability when designing and deploying AI systems. AI systems that interact with retail customers should be capable of identifying vulnerability signals and escalating to human support. An AI system that continues to process a consumer who is exhibiting vulnerability signals without human intervention may breach Consumer Duty requirements.

7.3 The regulatory perimeter must not be eroded by AI. TISA's response to the Mills Review identified that AI systems are already providing services functionally equivalent to regulated financial advice while remaining outside the regulatory perimeter. Firms should ensure that their own AI deployments do not cross the advice boundary without appropriate authorisation and controls. The FCA is expected to address the perimeter question in its Mills Review response.

7.4 Where AI agents are deployed to act autonomously on behalf of consumers, including in agentic finance management contexts, firms should apply standards equivalent to those applied to human advisers or discretionary managers. The degree of autonomy exercised by an AI agent does not reduce the firm's regulatory accountability for the outcomes it produces.

7.5 Deepfake and AI-enabled fraud represents a material and growing risk to consumers and to the integrity of firms' authentication processes. Firms should assess their vulnerability to AI-generated synthetic identities, voice cloning, and document forgery. Fraud defences should be reviewed against AI-enabled attack vectors on at least an annual basis.

7.6 Firms should consider the liquidity and market integrity risks that could arise from widespread AI-driven investment decision-making. If multiple AI systems trained on similar data make similar decisions simultaneously, the resultant market impact may be significant. This risk should be considered in the firm's risk framework and reported to the Board where material.

Example 7.1 Consumer Duty Assessment for an AI Chatbot

A retail bank deploys an AI chatbot to handle initial enquiries about mortgage products. The chatbot can answer questions about product features, eligibility criteria, and application processes.

SMF16 (Compliance Oversight) requires a Consumer Duty assessment before deployment. The assessment identifies three issues: the chatbot cannot identify vulnerability signals in written communication; the chatbot's responses to questions about affordability come close to personalised advice; and the chatbot does not provide consistent signposting to the bank's human advisers.

The chatbot is modified before deployment to include a vulnerability flagging protocol that escalates to a human agent when specific language patterns are detected, a defined script boundary that stops short of personalised recommendations, and clear and prominent links to human advice services.

SMF16 reviews outcome monitoring data quarterly. After three months, override rates (customers requesting a human agent) are tracked and consumer understanding scores from post-interaction surveys are reviewed.

Principle 8 Monitoring, Review and Continuous Improvement

AI governance frameworks and individual AI systems should be subject to ongoing monitoring, regular independent review, and continuous improvement. The pace of AI development requires that governance frameworks keep up with the technology they are designed to govern.

GUIDANCE

8.1 All AI systems should be monitored in production on a continuous basis. Monitoring should cover accuracy and performance metrics, rates of human override or escalation, data quality indicators and any incidents or near-misses. Monitoring data should be reported to the relevant SMF on at least a monthly basis for high-risk systems.

8.2 The firm's AI governance framework should be subject to formal review at least annually. The review should consider whether the framework remains appropriate in light of changes to the firm's AI use, changes to the regulatory environment and any material incidents involving AI systems in the firm or the wider industry.

8.3 Material incidents involving AI systems should be recorded in the firm's incident management framework. This includes incidents where AI systems produced incorrect outputs that affected customers or markets, incidents where AI systems were unavailable and caused operational disruption, and near misses where errors were caught before impact. Root cause analysis should be conducted for material incidents and findings should feed back into the risk assessment and governance framework.

8.4 Firms should actively engage with regulatory developments and industry working groups on AI governance. The FCA's AI Sprints, Sandboxes and Showcases provide opportunities to test approaches and engage with the regulator's thinking. Industry bodies including ISITC Europe provide forums for practical knowledge-sharing. Firms should contribute to and benefit from these resources.

8.5 Staff training and competency in AI oversight is a governance requirement, not merely a support function. Senior managers accountable for AI outcomes should receive training appropriate to their oversight responsibilities. This does not require technical AI expertise but does require familiarity with the risks AI presents, the questions they should be asking and the signs of inadequate control.

Example 8.1 Annual AI Governance Review

A UK investment manager conducts its annual AI governance review. The review is led by the SMF16 function with input from the COO function, the model risk team and the first line AI owners.

The review finds that two new AI systems have been deployed in the year without going through the full approval process. Both are classified as medium risk. The review recommends that the approval gateway be strengthened.

The review also notes that the FCA's Mills Review response, published earlier that year, includes new expectations on AI audit trail standards. The firm's audit trail requirements for two high-risk systems are assessed against the new standards and found to require enhancement.

The Board receives a summary of the review findings and approves a remediation plan with named owners and deadlines.

Annex A SMF Accountability Reference Table

The table below summarises the AI accountability allocation by SMF function. It is intended as a quick reference and should be read alongside Principle 3.

Firms that do not have a particular SMF function (for example, FCA solo-regulated firms that do not have SMF4 or SMF24) should allocate the relevant accountability to an equivalent executive function and document this clearly in the relevant Statements of Responsibilities.

SMF	Function	AI Accountability	Firm Type
SMF1	CEO	Ultimate accountability for AI strategy and risk appetite; sets ethical tone; regulatory accountability on AI matters	All firms
SMF2	Chief Finance Officer	Accountable for AI used in financial reporting, capital calculations, and stress testing; ensures AI-generated financial data meets audit standards	All firms where relevant
SMF3	Executive Director	In FCA solo firms: may hold risk, operational and technology AI accountability in absence of SMF4 and SMF24; must be specified in SoR	FCA solo firms
SMF4	Chief Risk Officer	Model risk management; independent model validation oversight; AI risk appetite and monitoring; bias and accuracy assurance	Dual-regulated firms only
SMF5	Head of Internal Audit	Independent assurance of AI governance framework; AI in scope of internal audit programme	All firms with IAF
SMF16	Compliance Oversight	Regulatory compliance of AI use cases; Consumer Duty assessment; oversight of regulatory perimeter; financial	All firms

SMF	Function	AI Accountability	Firm Type
		crime system compliance; SoR to reference AI oversight	
SMF17	MLRO	AML and financial crime AI systems; false positive management; suspicious activity reporting integrity; AI audit trail for financial crime	All firms with AML obligations
SMF24	Chief Operations Officer	Designated technical SMF; due diligence and validation function; operational resilience of AI; kill switch capability; CTP oversight of AI providers; SoR to reflect AI operational integrity	Dual-regulated firms only
Commissioning SMF	Varies	Business outcome accountability for AI deployed in their function; signs off risk classification; reliance on technical SMF documented	All firms

SMF4 and SMF24 are PRA-designated Senior Management Functions. They are applicable to banks, building societies, credit unions, PRA-regulated insurers, and other PRA-regulated entities. They do not apply to FCA solo-regulated firms. FCA solo-regulated firms must allocate equivalent accountability to their actual SMF structure and document this in Statements of Responsibilities.

SMF24 as the designated technical SMF reflects the function's natural ownership of operational integrity, technology, and vendor management. Where a firm has both an SMF24 and a Chief Information Officer or Chief Technology Officer who is not an SMF holder, the SMF24 should be supported by and may rely upon the expertise of the CIO or CTO function but retains regulatory accountability.

This table is indicative. Firms should take their own legal and regulatory advice on SMF allocation appropriate to their structure and activities.

Annex B Risk Classification Worked Examples

The examples below illustrate how the risk classification framework in Section 3 should be applied in practice. They are drawn from real use case categories discussed in industry forums including ISITC Europe and TISA.

B.1 Agentic AI for Investment Portfolio Management

An AI agent is deployed to manage discretionary investment portfolios for retail clients. The agent rebalances portfolios, executes trades and manages cash within parameters set by the client at account opening. The agent operates with minimal human intervention between quarterly reviews.

- Risk classification: High.
- Rationale: autonomous execution of financial decisions; retail client exposure; decisions may be difficult or costly to reverse; direct Consumer Duty implications; functionally equivalent to a discretionary investment manager
- Minimum controls: Board approval; SMF16 Consumer Duty assessment; independent model validation; clear customer disclosure of AI role; human override capability; full MiFID II-equivalent audit trail; quarterly human review of outcomes; SMF1 accountability for AI strategy; commissioning SMF accountability for outcomes

B.2 AI-Assisted Trade Surveillance

A market participant deploys an AI system to monitor trading activity for potential market abuse. The system generates alerts that are reviewed by a human surveillance team. No automated action is taken on AI alerts without human review.

- Risk classification: High (given market integrity implications)
- Rationale: market abuse detection is a core regulatory obligation; false negatives carry regulatory risk; human review is present but the AI filters what the humans see
- Minimum controls: named SMF16 accountability; independent validation of alert thresholds and false positive and negative rates; documented human review standard; audit trail linking each alert to its review outcome; annual review of model performance.

B.3 AI Summarisation Tool for Compliance Monitoring

A compliance team deploys an AI tool to summarise call recordings for quality assurance monitoring. The tool produces a summary and a risk score. A human compliance officer reviews the summary and score and decides whether to listen to the full call.

- Risk classification: Medium.
- Rationale: human decision-maker remains in the loop for all material outcomes; no direct customer-facing function; affects internal quality assurance only

- Minimum controls: documented human review requirement; accuracy monitoring against a sample of manually reviewed recordings; named business owner; data protection assessment for voice data processing; annual review.

B.4 Generative AI for Regulatory Intelligence Summarisation

A compliance function uses a generative AI tool to summarise regulatory publications and produce a briefing note for senior management. The tool operates on publicly available documents and produces outputs reviewed by a qualified compliance professional before distribution.

- Risk classification: Low to Medium (depending on whether output influences regulated decisions)
- Rationale: human expert review before any action is taken; output is information not instruction; no customer data involved
- Minimum controls: acceptable use policy; training for users on limitations of AI summarisation; human expert review requirement before distribution; periodic quality review

Annex C EU AI Act Crosswalk for Financial Services

The EU AI Act establishes a risk-based framework for AI systems. It classifies AI systems as unacceptable risk (prohibited), high risk, limited risk, or minimal risk. It also addresses general-purpose AI models, including large language models.

Financial services firms operating in the UK should understand how the Act applies to them and how its requirements relate to the Principles in this guidance.

C.1 High-Risk AI Systems under the EU AI Act

The EU AI Act classifies AI systems used in the following contexts as high risk where they affect access to financial services or creditworthiness assessment:

- creditworthiness assessment and credit scoring
- life and health insurance risk assessment and pricing (to the extent these affect access to insurance)

Additionally, AI systems used as safety components of critical infrastructure are classified as high risk. Firms should assess whether any AI systems embedded in payment infrastructure or financial market infrastructure fall within this category.

C.2 Obligations for High-Risk Systems

For high-risk AI systems, the EU AI Act requires:

- risk management systems to be established and maintained throughout the AI system lifecycle.
- data governance and management practices covering training, validation, and testing datasets.
- technical documentation sufficient to allow conformity assessment.
- record-keeping and logging to enable traceability.
- transparency and provision of information to deployers
- human oversight measures built into the system.
- appropriate levels of accuracy, robustness, and cybersecurity

These obligations map closely to the Principles in this guidance. Firms that comply with these Principles will be substantially prepared for EU AI Act compliance where it applies.

C.3 General-Purpose AI Models

The EU AI Act addresses general-purpose AI models, including large language models, directly. Providers of such models are subject to transparency, documentation, and safety obligations. Firms deploying such models as components of their AI systems should obtain the technical documentation required by the Act from their providers and incorporate it into their due diligence process (see Principle 4).

C.4 Divergence from UK Approach

The UK approach to AI regulation is currently less prescriptive than the EU AI Act. The UK Government has emphasised a principles-based, sector-led approach consistent with the FCA's existing outcomes-focused regulatory philosophy. The UK AI Bill, expected in May 2026, may introduce statutory high-risk categorisation but is not expected to replicate the full prescriptive requirements of the Act.

Firms operating in both the UK and EU, or serving EU customers, must assess which requirements apply to them and at what point. The Principles in this guidance are consistent with both the UK's current approach and the direction of travel of the EU AI Act. Complying with this guidance should not require significant additional effort to achieve EU AI Act compliance for the activities covered.

This crosswalk is provided for general information. The EU AI Act is a detailed and evolving instrument. Firms with significant EU operations or EU customer bases should take specialist legal advice on their obligations under the Act.