KPMG

CYBER RISK INSTITUTE

Deconstructing the Cyber Risk Institute FS AI RMF

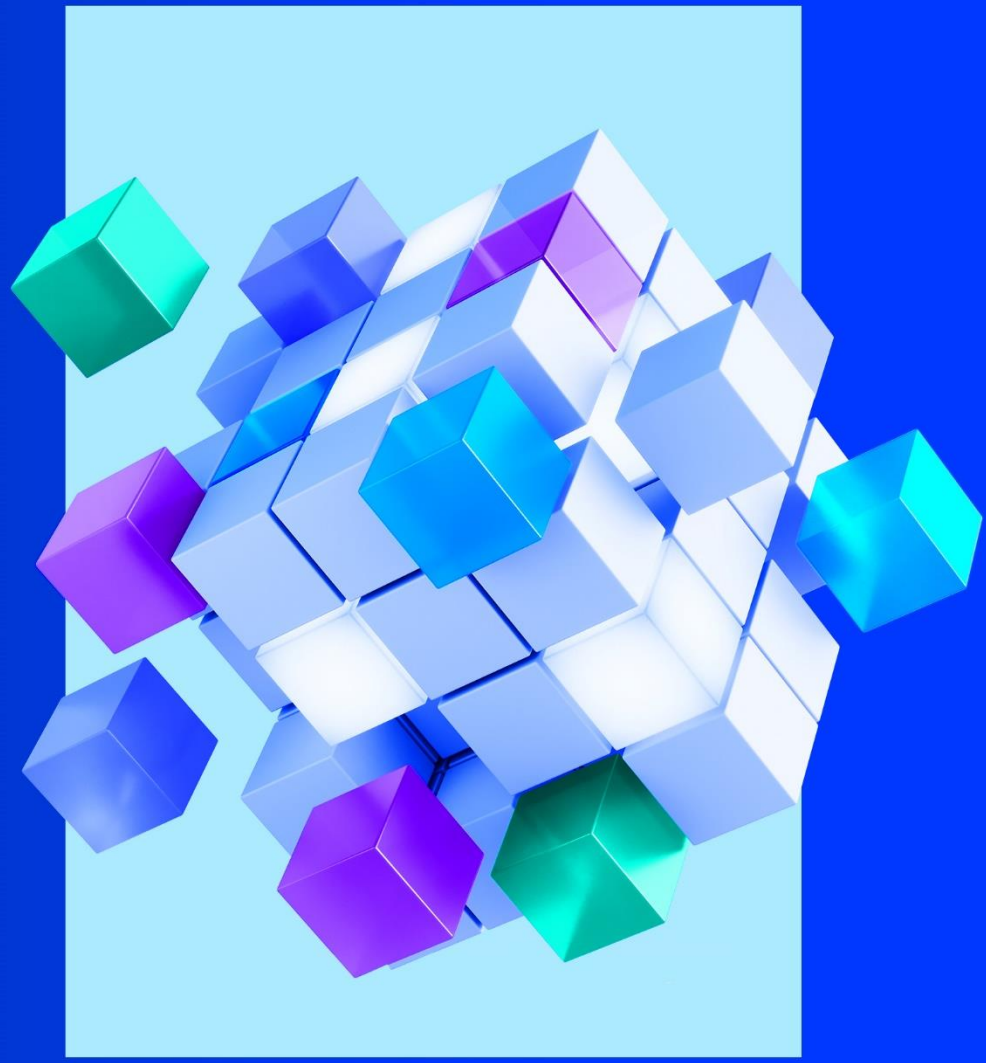

From KPMG contributors to the Cyber Risk Institute Financial Services Artificial Intelligence Risk Management Framework

Contents

About the CRI and the CRI FS AI RMF

About the Cyber Risk Institute

The Cyber Risk Institute (CRI) is a not-for-profit coalition of financial institutions and trade associations committed to advancing and unifying risk management standards for the financial sector, providing practical, industry-driven solutions for new and emerging risks in cybersecurity, technology, and artificial intelligence (AI).

CRI's commitment is demonstrated through a suite of three key solutions and frameworks: the foundational CRI Profile, its accompanying CRI Cloud Profile extension, and the recently published CRI Financial Services Artificial Intelligence Risk Management Framework ("CRI FS AI RMF" or "the Framework").[1]

About the CRI Financial Services Artificial Intelligence Risk Management Framework

The CRI Financial Services Artificial Intelligence Risk Management Framework (FS AI RMF) empowers financial organizations of all sizes with a common framework to identify, evaluate, manage, and govern the risks associated with AI, forged by industry leaders, and built upon broad industry consensus. The effort involved a diverse body of more than 100 financial entities, under the Financial Services Sector Coordinating Council (FSSCC) and its AI Executive Oversight Group, with participants ranging from community banks and credit unions to national and multinational banks, investment firms, insurance companies, and trade associations. U.S. and international agencies, most notably the National Institute of Standards and Technology (NIST), also played a pivotal role.

Alignment with existing frameworks, standards, and regulatory guidance

The CRI FS AI RMF is engineered to integrate with and enhance an organization's existing risk management functions, serving as a universal supplement – not a replacement. It does not function as an exhaustive catalog of AI risks or controls, nor does it prescribe role-specific instructions. Instead, by treating AI as an enterprise-wide risk, the Framework is built to connect logically with broader governance, risk, and compliance (GRC) programs, thereby reinforcing an organization's integrated risk assessment,[2] aggregation, and mitigation strategies.

Alignment with existing risk management functions

Grounded in the NIST AI Risk Management Framework (NIST AI RMF), the CRI FS AI RMF adapts and extends its structure for the financial sector, synthesizing global regulations, standards, and supervisory provisions **(Figure 1)**. This extension augments NIST's pre-existing Functions, Categories, and Sub-Categories, introducing 230 distinct Control Objectives with the goal of contextualizing, not replacing, standards and promoting global harmonization.

Figure 1: Select Primary and Secondary Informative References for the CRI FS AI RMF. The framework aligns with and is informed by a broad set of key global standards, regulations, and principles.

- NIST AI 600-1 (AI RMF 1.0)
- SR 11-7: Supervisory Guidance on Model Risk Management
- FFIEC Guidance
- EU AI Act
- CRI Profile
- ISO/IEC 42001: Information technology – Artificial intelligence – Management system
- MAS Guidelines of AI Risk Management for Financial Institutions
- OECD G7 Toolkit for AI in the Public Sector

[1] The CRI Profile, Cloud Profile, and FS AI RMF are available for download at no cost on the CRI website: https://cyberriskinstitute.org/the-profile/.
[2] The CRI FS AI RMF itself is not considered a risk assessment.

CRI FS AI RMF background, purpose, and intended audience

Background and motivation

The rapid adoption of AI in the financial sector presents a dual reality: while offering significant opportunities (e.g., enhanced operational efficiency, improved fraud detection, and greater personalization), it also introduces novel risks (e.g., algorithmic bias, lack of transparency, and cybersecurity vulnerabilities). In addition, although robust for conventional technologies, traditional risk management frameworks were not designed to mitigate the novel challenges posed by advanced AI. Additionally, many organizations believe they align with the defined NIST AI RMF principles, but applying these principles to advanced AI systems, such as large language models (LLMs), reveals gaps due to complexity and novelty.[3]

As a result, industry experts are emphasizing that effectively governing these advanced models necessitates more practical and targeted guidance than conventional frameworks provide. Therefore, the industry desires a tailored AI risk management approach, one that is designed to keep pace with the evolution of AI and supports financial entities in deploying trustworthy and responsible AI.

Purpose and benefits

The CRI FS AI RMF provides the financial sector with a common, practical standard for managing AI risks. It is designed not as a replacement for existing standards, but as a supplement that adapts and extends the NIST AI RMF, contextualizing its principles for financial services.

The CRI FS AI RMF value proposition:

- **Strengthens AI governance**: Provides a robust and structured approach to identifying, managing, and governing AI-specific risks.
- **Fosters trust**: Builds confidence among regulators, customers, and internal stakeholders by demonstrating a focus on responsible AI.
- **Enables responsible innovation**: Offers a clear pathway for deploying AI innovations securely and transparently, allowing firms to realize benefits while managing risks.
- **Provides actionable clarity**: Bridges the gap between high-level principles and technical reality with 230 distinct Control Objectives.
- **Promotes global harmonization**: Establishes a common language and standard for managing AI risk, facilitating greater consistency across the global financial ecosystem.

Intended audience

The CRI FS AI RMF is designed for broad applicability across the financial industry, intended for use by financial institutions of any size, complexity, or type, as well as their third-party providers throughout the AI supply chain.

While the Framework is role-agnostic, it is valuable for a diverse range of roles, including AI professionals (e.g., engineers, data scientists, architects, developers, researchers), enterprise technology leaders, risk and compliance managers, and legal professionals. The Framework is also highly relevant for professionals at the intersection of finance and technology, particularly those with deep expertise in financial markets, regulations, and AI/ML systems. To effectively apply the Framework, users require a baseline understanding of both AI and associated risk management principles.

[3] U.S. Department of the Treasury Report on the Uses, Opportunities, and Risks of Artificial Intelligence in Financial Services: https://home.treasury.gov/news/press-releases/jy2760

Deconstructing the CRI FS AI RMF into its core building blocks

The CRI FS AI RMF has four distinct yet integrated components: AI Adoption Stage Questionnaire (“Questionnaire”), Risk and Control Matrix (“RCM”), Guidebook, and Control Objective Reference Guide **(Figure 2)**.

These components are designed to work in concert, informing an organization’s approach to AI risk management. When using the CRI FS AI RMF, an organization begins with the Questionnaire to assess its current AI adoption. The results then inform which elements within the RCM are most relevant, allowing the user to focus on the appropriate risks and controls for their specific stage. The Guidebook provides overarching instructions for how to leverage and deploy the CRI FS AI RMF, while the Control Objective Reference Guide serves as a detailed supplement for each Control Objective in the RCM.

Figure 2: The CRI FS AI RMF's building blocks. This figure highlights the four primary components of the framework: the AI Adoption Stage Questionnaire, the Risk and Control Matrix, the Guidebook, and the Control Objective Reference Guide.

The **AI Adoption Stage Questionnaire** serves as the entry point to the Framework. It guides a user to determine their organization's AI adoption by having them match their current operational practices to a series of stage-specific descriptions. The resulting classification is essential for leveraging the remainder of the Framework effectively.

The **Guidebook** serves as the primary user manual for the entire Framework. It provides step-by-step instructions and best practices for using the Questionnaire and the RCM effectively and for integrating the Framework into an organization's existing risk management strategy.

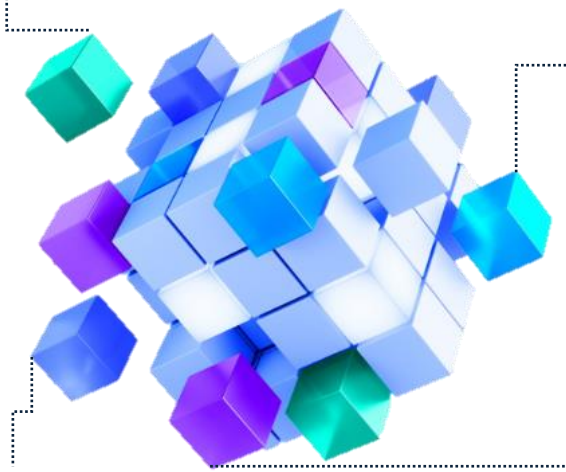

The **Risk and Control Matrix** is the core component of the Framework, presenting a comprehensive matrix of potential Risk Statements mapped to their corresponding Control Objectives. Additionally, the RCM offers practical implementation guidance that is specifically tailored to each Adoption Stage, providing recommendations that are relevant to the organization’s adoption level.

The **Control Objective Reference Guide** is a supplementary resource that provides a comprehensive breakdown of each Control Objective found in the RCM, detailing the Control Objective's definition and scope and providing practical examples of corresponding controls and evidence to demonstrate their effectiveness.

Understand your current state with the AI Adoption Stage Questionnaire

The AI Adoption Stage Questionnaire guides organizations through a self-assessment designed to pinpoint its current AI Adoption Stage. This classification places the organization into one of four stages of adoption–Initial, Minimal, Evolving, or Embedded–determined by three key considerations: business impact of AI, technology implementation, and scalability **(Figure 3)**.

To determine their organization's appropriate AI Adoption Stage, the Questionnaire prompts the user to evaluate six key dimensions–Business Impact, Governance, Deployment Model, Third-Party AI Use, Organizational Goals, and Data Sensitivity and Criticality–by comparing their current practices and policies against statements that reflect varying operational states.

Based on the response, the outcome is a clear classification of the organization into an early, mid, or late stage of AI adoption, enabling a targeted application of the RCM.

Figure 3: AI Adoption Progression. This figure illustrates the journey of AI adoption, from the Initial to the Embedded stage. It defines the progression of adoption for each stage across the three key considerations of business impact, technology implementation, and scalability.

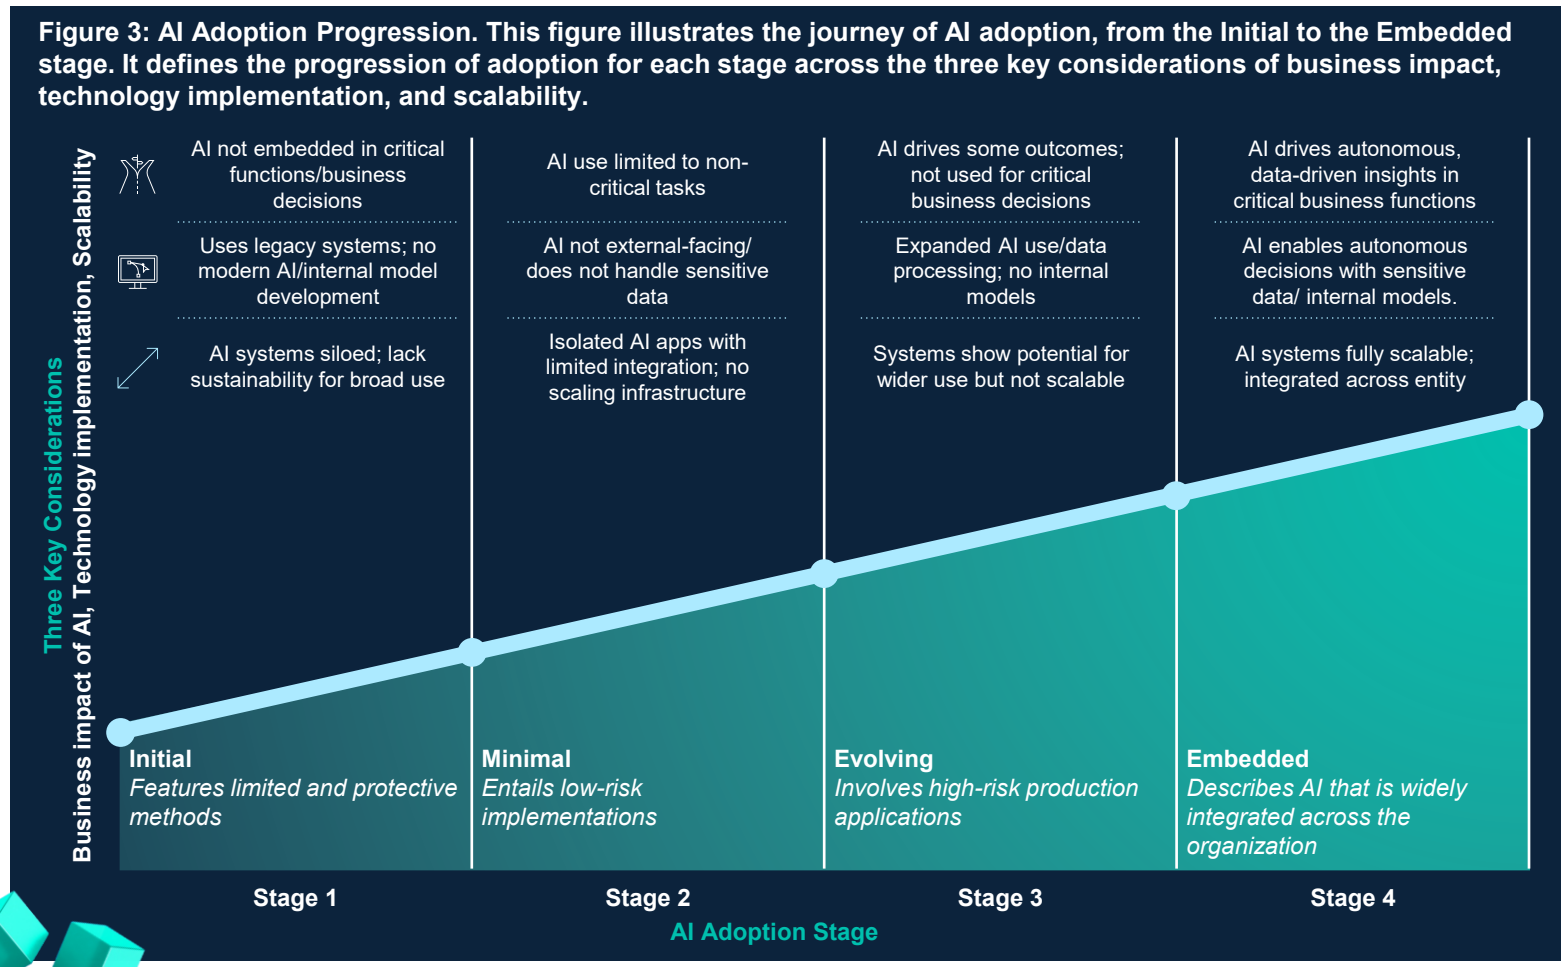

Translate AI adoption into prioritized risks and controls with the Risk and Control Matrix

The RCM is the central component of the Framework, providing a comprehensive and structured view of each Control Objective and its associated information. For each Control Objective, the RCM defines the relevant AI Adoption Stage(s), provides implementation guidance, and maps it to the associated Risk Statement, the corresponding AI Trustworthy Principle, and the relevant NIST AI RMF hierarchy (Function, Category, and Sub-Category). A key feature of the RCM is that the Control Objectives are aligned with the Adoption Stage(s) which they apply. For example, a Control Objective may only apply to organizations at the Evolving or Embedded stage. This enables early-stage organizations, at the Initial or Minimal stages, to focus their resources effectively, concentrating on the most pertinent controls while deferring those intended for later stages of adoption.

NIST Function, Category, and Sub-Category

To align with industry standards, the RCM adopts the hierarchy of the NIST AI RMF, organizing its structure around the established Functions (i.e., Govern, Map, Measure, and Manage), Categories, and Sub-Categories. Each includes a Reference ID, Name, and Description. Each Control Objective is aligned to the relevant Sub-Category, which is nested within a broader Category, falling under an overarching Function. This classification provides a standardized hierarchy for organizing risk and control information, promoting consistency and clarity in risk management practices.

AI Trustworthy Principle and Risk Statement

Each Control Objective is mapped to its primary NIST AI Trustworthy Principle and a Risk Statement, allowing for prioritization and aggregation within the Framework. In practice, multiple AI Trustworthy Principles may apply, and Risk Statements may be tailored to align with a firm's internal risk .

The CRI FS AI RMF is grounded in the NIST AI Trustworthy Principles which include safe, secure and resilient, explainable and interpretable, privacy-enhanced, fair, valid and reliable, and accountable and transparent.

Control Objectives

Control Objectives are the actionable components within each Sub-Category, representing the distinct policies, processes, or operational steps to mitigate AI-related risks. Each includes a Reference ID, Name, and Description. Control Objectives span a wide range of topics including data quality, fairness, operational resiliency, security, and transparency.

The RCM currently contains 230 Control Objectives. However, the number of applicable Control Objectives increases with each Adoption Stage, starting at 21 (Initial) to all 230 (Embedded). Across the Functions, Govern includes 81, Map includes 47, Measure includes 59, and Manage includes 43.

Implementation Guidance

The Implementation Guidance offers practical recommendations for applying each Control Objective. It is not intended as an exhaustive set of instructions, rather as a directional starting point to help organizations translate the Control Objectives into specific actions that suit their unique needs and environments.

The Risk and Control Matrix structure and organization

The RCM is intentionally dynamic, allowing the information it contains to be viewed from various perspectives.

However, its underlying structure is grounded in NIST's Functions, Categories, and Sub-Categories, with CRI-defined Control Objectives organized at the Sub-Category level.

Each of these components has an ID, a Name, and a Description **(Figure 4).** Each Sub-Category generally contains between three and five Control Objectives, although this number can vary.

Figure 4: An Illustration of the RCM's Hierarchical Structure. This figure provides a practical example of the RCM's structure. It demonstrates how a specific Control Objective is nested within a Sub-Category, which in turn rolls up into a Category under an overarching Function, with each component defined by a unique ID, Name, and Description.

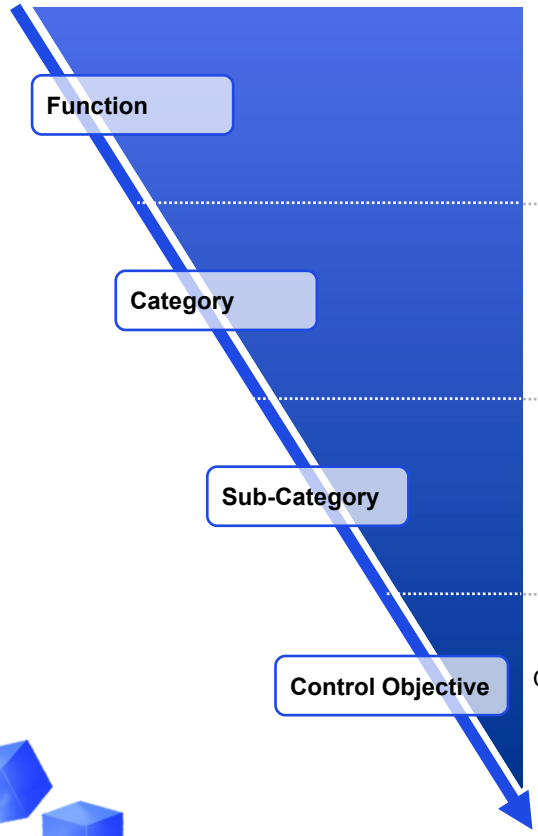

	ID	Name	Description
Function	GV	Govern	A culture of risk management is cultivated and present.
Category	GV-1	Establishing Key Policies & Processes	Policies, processes, procedures, and practices across the organization related to the mapping, measuring, and managing of AI risks are in place, transparent, and implemented effectively.
Sub-Category	GV-1.1	Legal & Regulatory Compliance	Legal and regulatory requirements involving AI are understood, managed, and documented.
Control Objective	GV-1.1.1	AI Legal, Regulatory, & Policy Integration	The organization identifies, monitors, and integrates applicable laws, regulations, contractual obligations, and sector requirements into policies, procedures, and operations governing AI. This includes updating governance artifacts, aligning operational practices, and verifying ongoing compliance as requirements evolve.

Operationalize AI risk management with the Guidebook

The Guidebook aims to promote sector-wide adoption and understanding of the Framework and help organizations effectively implement the Framework. It provides detailed guidance on the 230 Control Objectives captured in the RCM, aligned with the four stages of AI adoption defined in the AI Adoption Stage Questionnaire. The Guidebook reflects sector-specific considerations and global standards, guidelines, and regulatory expectations to support consistent evaluation, benchmarking, and adoption in AI governance. It offers practical guidance for developing, implementing, and maturing AI risk and control programs, fostering trust among stakeholders, and demonstrating a commitment to responsible AI deployment.

The Guidebook's companion document, "Control Objective Reference Guide," offers additional details, including examples of Controls and Effective Evidence. Both documents support organizations in understanding how to interpret and apply this information at various stages of AI adoption, from initial AI deployment to deeply embedded, enterprise-wide AI implementation.

The Guidebook includes:

- **Section 1:** The AI Adoption Stage Questionnaire and instructions for completing it
- **Section 2:** An overview of the RCM and its contents
- **Appendix A:** NIST AI Trustworthy Principles
- **Appendix B:** Acronyms and Abbreviations
- **Appendix C:** Glossary of Key Terms
- **Appendix D:** Informative References
- **Appendix E:** Additional Source Materials

and enable control design with the Reference Guide

The Control Objective Reference Guide serves as a companion to the main Guidebook, providing a deeper layer of practical detail. For each of the 230 Control Objectives, it offers five illustrative controls, each with a clear description of the action, complemented by four examples of effective evidence to demonstrate successful implementation and adherence.

The example Controls and Effective Evidence in the Control Objective Reference Guide are illustrative, not exhaustive. They are intended to serve as directional guidance and not to be interpreted as a definitive guarantee of meeting regulatory or auditing expectations, as their suitability will vary based on each organization's unique context. Therefore, users must critically assess the relevance and adequacy of the provided examples for their own environment. Organizations should weigh the value of any evidence against the resources required to produce it and are encouraged to tailor their approach.

Getting started with the CRI FS AI RMF

Adopting the CRI FS AI RMF can be achieved through a high-level, five-step action plan, beginning with convening a cross-functional working group and completing the Questionnaire. Once completing the Questionnaire, a gap analysis may be performed against the RCM to identify opportunities for improvement. These findings can then inform the development of a prioritized implementation roadmap, with the Guidebook and Control Objective Reference Guide serving as key resources.

01

Assemble a cross-function working group

Form a working group of key stakeholders from relevant departments, if not yet established, appointing from departments such as Technology (e.g., AI Developers), Risk and Compliance (e.g., Model Risk and Legal), Business Lines (e.g., Product Owners using AI), and Executive Leadership (e.g., Sponsor/ Champion).

02

Complete the Questionnaire

Leverage the AI Adoption Stage Questionnaire to conduct a self-assessment of your current practices, which will provide a clear classification of your organization into one of the four stages of AI adoption.

03

Apply the RCM

Benchmark your current AI governance practices against the Control Objectives for your current and desired adoption stages to pinpoint strengths and opportunities for improvement.

04

Develop a roadmap

Develop a strategic and prioritized implementation roadmap that outlines the sequence and timing for adopting relevant Control Objectives, enabling your organization to systematically bridge existing gaps and mature its AI risk management capabilities over time.

05

Leverage the Guidebook and the Control Objective Reference Guide

Use the detailed guidance, examples, and references to support implementation.

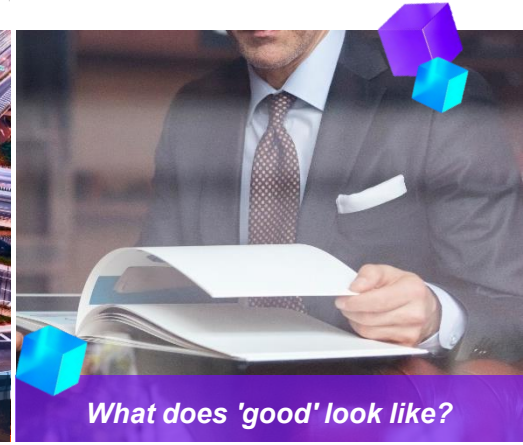

How KPMG can help

Work with authors behind the Framework

As key contributors to the CRI FS AI RMF, KPMG offers distinct insight and experience to help your organization navigate its adoption and implementation. Our team possess a first-hand understanding of the Framework's principles, components, and intended application. This distinct position allows us to move beyond theoretical guidance and provide practical, actionable guidance tailored to the specific challenges and opportunities within financial services. Our multidisciplinary teams can help your organization effectively and efficiently implement the CRI FS AI RMF, regardless of your current AI Adoption Stage.

Our services are designed to provide a wide array of support for your AI risk management program:

AI adoption and strategy

Our professionals facilitate the initial adoption process, helping you form a cross-functional working group, establish governance capabilities, and develop a clear strategy for leveraging AI in alignment with your business goals and regulatory requirements.

Current state assessment

We perform a current state assessment of your AI program using the AI Adoption Stage Questionnaire to benchmark your current capabilities and identify your Adoption Stage. We then conduct a gap analysis of your existing practices against the CRI FS AI RMF RCM Control Objectives to identify strengths and opportunities for improvement.

Based on the gap analysis results, we collaborate with your team to develop a prioritized, actionable roadmap for implementing the necessary controls, maturing your AI risk management capabilities.

Control implementation

We provide direct support to design and implement new AI controls, policies, and processes, promoting effective integration into your existing governance, risk, and compliance structures.

Contact us

Matt Miller
Principal, Cybersecurity & Tech Risk
T: +1 347 638 4726
E: matthewpmiller@KPMG.com

Karolina Oseckyte
Director, Cybersecurity & Tech Risk
T: +1 305 304 9299
E: karolinaoseckyte@KPMG.com

Breah Sandoval
Director, Cybersecurity & Tech Risk
T: +1 227 255 1644
E: breahsandoval@kpmg.com

Joe Crouse
Director, Cybersecurity & Tech Risk
T: +1 704 930 8524
E: josephcrouse@KPMG.com

Contributors

Ashley Ryan
Senior Associate, Cybersecurity & Tech Risk
T: +1 404 652 0405
E: ashleyryan@KPMG.com

Learn about us: **kpmg.com**